BEST
MINI PCs PG.12

REVIEW
RTX 5080 PG.74

DETECT AI
CONTENT PG.60

MAXIMUMPC

MINIMUM BS • APRIL 2025 • www.maximumpc.com

BUILD A 5080 MICRO PC!

- Craft this games console-sized PC
- Next-gen RTX GPU in an SFF case
- PS5 Pro-killing performance

BUILD THIS PC
STEP-BY-STEP
GUIDE
PG.16

NVIDIA
BLACKWELL
DEEP DIVE PG. 32

FUTURE
VOL 30, NO 4

THE RISE AND
FALL OF
PACKARD BELL PG. 44

PC GAMER
EDITOR'S CHOICE
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM
Digital Storm PCs featuring Intel® Core™ i7 processors.

intel
CORE i7
inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

inside

APRIL 2025

SCAN TO GET THE
TOM'S HARDWARE
WEEKLY NEWSLETTER

QUICKSTART

8 **THE NEWS**
AMD announces Radeon RX 9000; Microsoft rolls out Windows 11; Twitch monetization changes.

12 **THE LIST**
The best mini PCs.

If you're after a tiny, off-the-shelf PC with awesome everyday computing power.

13 **TRADE CHAT**
Tariffs are here—and they'll have consequences for prices.

48 **CENTERFOLD**
Corsair Vengeance A7500.

R&D

57 **HOW TO**
Detect AI content; Build your own streaming server; Beat games with Edge Game Assist; Pass on your passwords.

LETTERS

14 **DOCTOR**

94 **COMMENTS**

16 **BUILD AN RTX 5080 MICRO PC!**

32 **NVIDIA BLACKWELL**
We explore Nvidia's newest chips, and how AI has changed the game.

44 **PACKARD BELL**
Charting the rise and fall of the once-ubiquitous PC brand.

46 **ITX: A HISTORY**
How this changed everything about PC, and continues to have an impact.

52 **BUILD A STEAM DECK**
Fancy the ultimate handheld gaming PC? *John Knight* shows you how.

© 2K GAMES, BEE-LINK

IN THE LAB

74 **NVIDIA GEFORCE RTX 5080**

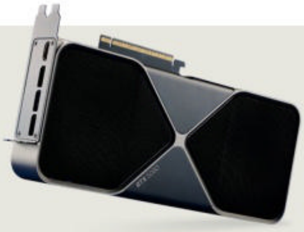

76 **DOUGH SPECTRUM BLACK 32 OLED**

86 **GIGABYTE X870 AORUS ELITE WIFI7 ICE**

88 **LOGITECH G RS WHEEL HUB & TRACK WHEEL**

90 **CIVILIZATION 7**

AMAZING
SUBSCRIPTION OFFERS
SUBSCRIBE TO ANY OF OUR **BEST-SELLING** MAGAZINES

RUGBY WORLD
YOUR 2021 LIONS

FourFourTwo
EURO 2020
THE WAIT IS OVER

PLAY
BATTLEFIELD 2042

retro GAMER
THE STAR WARS SAGA

Total Guitar
BACK TO LIVE!

GOLF MONTHLY
Beat your best score

SHOP THE FULL RANGE
WWW.MAGAZINESDIRECT.COM/**SUBSCRIBE**

Offer open to new subscribers only. Please allow up to six weeks for delivery of your first subscription issue. Payment is non-refundable after the 14 day cancellation period unless exceptional circumstances apply. For full terms and conditions, visit www.magazinesdirect.com/terms. For enquiries please call: +44 (0) 330 333 1113. Lines are open weekdays between 9am-5pm UK Time or e-mail: help@magazinesdirect.com. Calls to 0330 numbers will be charged at no more than a national landline call, and may be included in your phone provider's call bundle.

Meet the technology experts

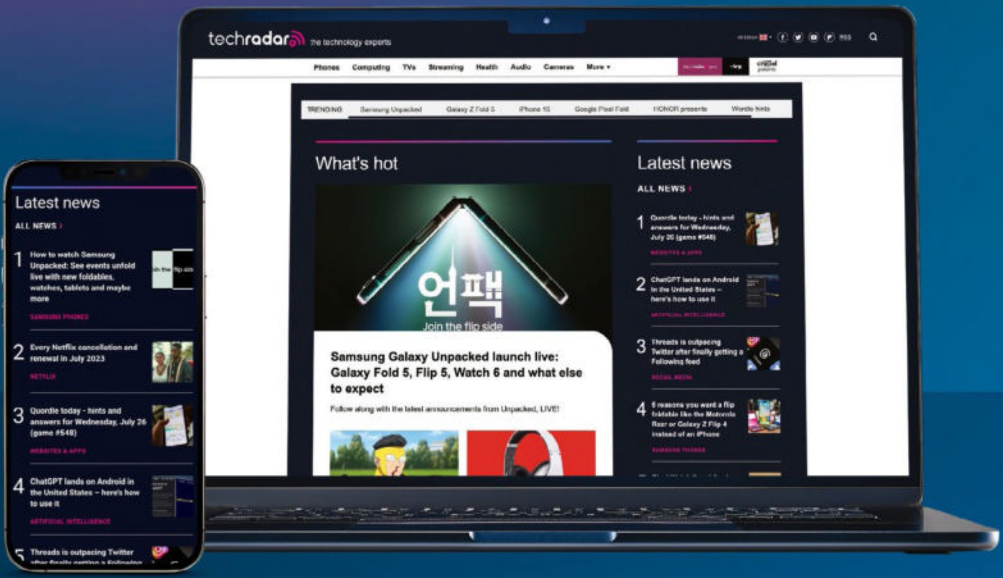

techradar the technology experts
Latest news
What's hot
Join the flip side

Samsung Galaxy Unpacked launch live: Galaxy Fold 5, Flip 5, Watch 6 and what else to expect

Follow along with the latest announcements from Unpacked, LIVE!

- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Contributing Writers: Reece Bithrey, Nate Drake, Ian Evenden, Dave James, John Knight, Jeremy Laird, Anthony Leather, Nick Peers, Nik Rawlinson, Madeline Ricchiuto, Jacob Ridley, Zak Storey, Jarred Walton, Robert Zak
Production Editor: Steve Wright
Editors Emeritus: Andrew Sanchez, Gordon Mah Ung

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: geekompc.com, Future PLC

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Mike Peralta, mike.peralta@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com

MANAGEMENT
CEO: Jon Steinberg
Brand Director: Faisal Alani
Content Director: Marc Chacksfield
Group Art Director: Warren Brown

PRODUCTION
Head of Production: Mark Constance
Senior Production Manager: Matthew Eglinton
Production Manager: Vivienne Calvert
Production Assistant: Emily Wood

Future US LLC, 130 West 42nd Street,
7th Floor, New York. NY 10036. USA. www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION
Maximum PC is available for licensing and syndication.
To find out more, contact us at licensing@futurenet.com or view our available content at www.futurecontenthub.com.
Head of Print Licensing: Rachel Shaw, licensing@futurenet.com

SUBSCRIBER CUSTOMER SERVICE
Website: www.magazinesdirect.com
Tel: 844-779-2822
New Orders: help@magazinesdirect.com
Customer Service: help@mymagazine.co.uk

BACK ISSUES
Website: https://bit.ly/mpcsingleissue

Next Issue On Sale April 22, 2025

EXTRA DIGITAL FEATURES

AUDIO FILE | PHOTO GALLERY | VIDEO FILE

© 2025 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors.
Creators.
Experience
Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief Executive Officer **Jon Steinberg**
Non-Executive Chairman **Richard Huntingford**
Chief Financial Officer **Sharjeel Suleman**

Tel +44 (0)1225 442 244

© GETTY IMAGES

Guy Cocker

HONEY, I SHRUNK THE PC!

IF THERE'S ONE THING you can rely on in technology, it's the ever-onward march to miniaturize anything with a power button. Looking at Nvidia's 40-series graphics card range, you'd think Jensen Huang had missed the memo. But along come the 50-series cards, and compared to the chunky 40-series, they're positively titchy. This can mean only one thing—it's small form factor build season here at *Maximum PC*.

We're taking a hammer to the laws of physics to squeeze a meaty, full-fat RTX 5080 GPU into a case the size of a box of cereal in our cover build (page 16). We're talking about a console-conquering micro PC that could sit on your desk without scrimping on performance or cooling. Our guide will help you pick all the components, from the case through the CPU and beyond, and take the headaches out of keeping cables in check, optimizing air flow, and all the other fiddly bits. You'll be the envy of PS5 peasantry and Xbox blowhards.

If getting down and dirty with cable management and fan curves gives you the chills, fear not—we've also chosen our favorite pre-built micro PCs on page 12, which all boast oodles of power, despite being not much bigger than a Rubik's cube.

We also put that Nvidia RTX 5080 under the microscope in our lab test this issue. Half the price of the top-tier 5090, it's the sensible(ish) option for anyone that skipped the 40-series generation, and brings some smart AI features to the table. Multi Frame Generation is here to stay, and thankfully, its more-frames cheat code doesn't result in a latency gulf or unsightly visual smearing. You can read more about how Nvidia's managed all this in Jarred Walton's Blackwell GPU deep dive, which covers what makes these AI-fueled graphics cards tick.

Keep being duped by AI content? The rise of generative artificial intelligence models is making it easier than ever to knock up convincing fakes of everything from a CV to a Hollywood blockbuster. We'll show you how to separate AI content from the carefully crafted work of human hands—and hopefully delay a *Terminator*-style AI apocalypse for a few more months, which you can read on page 60.

It's a far cry from the landscape that Packard Bell inhabited in its heyday, then—an LA computing giant that saw its fortunes fall apart with the turn of the millennium. In a new regular feature, we look back on the now-forgotten brand, and see what's left of a former industry mainstay, on page 44.

Elsewhere, we've got the return of a classic PC gaming franchise. *Civilization 7* is here to squeeze 'just one more turn out' of our digital world-domination plans. It's a corker, with our reviewer characterizing it as much as an all-new take on 4X strategy as it is a sequel—but does it better *Civ 6*?

If you like your gaming sessions with more gearboxes than Gandhis, the Logitech G RS Wheel Hub and RS Track Wheel will have you tearing up the tarmac from the comfort of your home. It's up there with the best racing wheel setups we've tested—easy to install, great in the hands, and a fair sight cheaper than a day out on the test track.

The theme may be micro formats, then, but it's another bumper issue. Grab those shrink rays and dive on in!

Guy Cocker

Guy is Maximum PC's *editor-in-chief. He built his first gaming PC in 1997 to play* Tomb Raider *on 3dfx, and has been obsessed with all things PC ever since.*

↘ submit your questions to: editor@maximumpc.com

THE NEWS

AMD Announces Radeon RX 9000

Next-gen GPUs promise significant performance gains

AMD HAS officially announced the RX 9000 series, starting with the Radeon RX 9070 and RX 9070 XT. The series and RDNA4 architecture were partially revealed at CES 2025 during AMD's keynote address. However the company didn't reveal in-depth technical information, partly due to the announcement taking place before Nvidia's address concerning its GeForce RTX 50-series GPUs.

The specifications and pricing for the RX 9000 series were the subject of rumors and leaks, specifically that AMD had delayed its rollout. The company has now ended all speculation by releasing full technical specifications for both GPUs, along with its own performance benchmarks.

Both cards feature 16GB of GDDR6 memory on a 256-bit interface, running at 20 Gbps. The AMD Radeon RX 9070 XT has 64 compute units, while the RX 9070 has 56. In terms of game clock, AMD puts the RX 9070 at 2.1GHz and RX 9070 XT at 2.4GHz. The boost clock is also up to 2.5GHz and 3.0GHz. The flagship XT has a total board power (TBP) of 304W, while the RX 9070 draws 220W.

At the time of writing, we couldn't find in-depth independent benchmarks. However, AMD asserts that the RX 9070 XT delivers over 40 percent superior performance at 1440p resolution compared to the previous generation AMD RX 7900 GRE. AMD also claims that the RX 9070 offers a 20 percent improvement over the same card. At 4K, the 9070 XT delivered 37 percent better performance in rasterization mode, while the 9070 outperformed its predecessor by 18 percent.

For ray traced games, the 9070 XT was 53 percent and 49 percent faster than the RX 7900 GRE at 4K and 1440p. The RX 9070 was 22 percent and 20 percent faster. These tests were performed by AMD Performance Labs on Windows 11 PCs using 43 different titles.

In terms of raw computing power, the 9070 XT works out at 48.7 TFLOPS FP32 and the 9070 at 36.1 TFLOPS. In theory, this makes the XT chip around one third faster, though real world performance may differ. AMD has also upgraded the AI accelerators in the compute units—each can now do twice as many FP16 operations. They now support sparse operations, potentially doubling performance.

AMD has also been working on the latest version of its AI-enhanced 'Super Resolution Technology', FSR4. This is currently exclusive to RDNA 4 GPUs. It leverages the AI accelerators to deliver superior image quality compared to FSR 3.1

AMD's flagship GPU boasts 64 CUs and 16GB of GDDR6 memory.

© AMD

At present, AMD has promised FSR 4 support for over 30 games. While support for more is anticipated, this places AMD far behind Nvidia. The rival DLSS 4 upscaling technology was compatible with 75 games at launch. This can be boosted further through the DLSS override feature, which swaps a game's DLL files for the latest DLSS versions, even if the title isn't officially supported.

The new RDNA 4 architecture incorporates a dual media engine for enhanced game streaming and recording. AMD claims this offers a 25 percent gain in H.264 low-latency quality, as well as an 11 percent improvement in HEVC encode quality. The overall encoding performance boost at 720p is estimated at 30 percent. AMD also claims a 50 percent performance decode uplift for AV1 and VP9.

RDNA 4 also introduces 'out of order' memory boxes. With RDNA 3, memory requests from different shaders were handled in a certain sequence. This meant that even when data items were returned quickly, they could be held up behind long-latency requests. The latest RDNA 4 architecture allows requests to be satisfied in any order. This lets shaders execute efficiently, regardless of long latency requests.

With RDNA 4, shaders can also allocate live registers dynamically. Registers can now be requested from the available pool, then released when the work is done. This allows for better handling of memory latency.

The MSRP for the AMD Radeon RX 9070 XT is $599, while pricing for the RX 9070 starts at $549. **–ND**

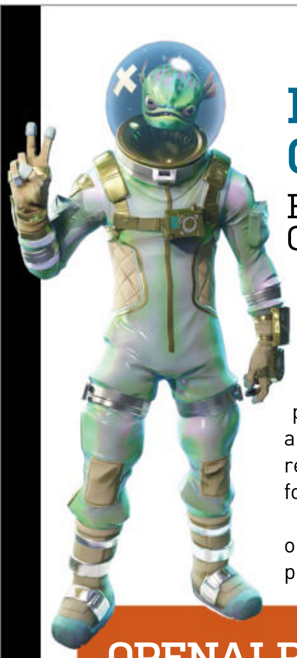

FORTNITE CHEATERS PROSPER

EPIC OFFERS CHEATERS A SECOND CHANCE WITH NEW BAN POLICY

IN A DEPARTURE from its previous zero-tolerance policy towards cheating, Epic Games has announced that it will offer a 'second chance' to certain *Fortnite* players.

From April onward, cheating will be punished with a one-year matchmaking ban. Epic Games states that this will allow people 'to learn from their mistakes', although those who cheat a second time will be banned for life. This new policy will also be retrospective, so players who have been subject to a lifetime ban for over one year will have their access to *Fortnite* restored.

The rules haven't changed for players who have sold cheats or engaged in other behaviors that have resulted in a lifetime ban previously. Such accounts will still be barred permanently. **–ND**

OPENAI PREVIEWS GPT-4.5

The latest model features more intuitive chat and fewer hallucinations.

OPENAI HAS released a 'research preview' of its flagship language model, which it describes as its 'largest and best for chat yet'. Anyone willing to pay $200 per month for a 'Chat Pro' subscription can try out version 4.5 today.

This 'non-reasoning' model, known internally as 'Orion', is designed to better interpret user intent and have a higher virtual 'EQ' when it comes to interactions. Nevertheless, the company has emphasised that GPT-4.5 is not one of its cutting-edge 'frontier' models. Even so, it claims that the LLM hallucinates less than GPT-4o, and is better at managing conversations.

OpenAI also tested GPT-4.5 using its 'SimpleQA' knowledge quiz. The model scored 62.5 percent. This compares well with GPT-4o (38.6 percent) and o3-mini (15 percent).

CEO Sam Altman tweeted that free tier users will benefit from unlimited chat access. However, there's currently no official release date. **–ND**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

✓ **SIGNAL THREATENS TO LEAVE SWEDEN**
They will not cooperate with proposals by the government for backdoor encryption products.

✓ **GTA V FOR PC UPDATES**
The Enhanced edition adds support for upscaling via Nvidia DLSS 2 and AMD FSR1/FSR3.

✓ **RAZER RELEASES MINECRAFT ACCESSORIES**
The gaming brand is 'Creepin' it real' with a Minecraft edition gaming mouse and headset.

TRAGEDIES

✗ **PASSWORDS STOLEN**
Troy Hunt, the owner of 'Have I Been Pwned?', has confirmed that the 'ALIEN TXTBASE' contains stolen credentials.

✗ **UBLOCK ORIGIN DEAD**
'uBlock Lite' lives on in Chrome, and uBlock Origin works in non Chromium-based browsers.

✗ **APPLE DISABLES ICLOUD ENCRYPTION IN UK**
Apple has disabled the service for UK users rather than comply with the government.

© MICROSOFT, EPIC GAMES, OPENAI

MICROSOFT ROLLS OUT WINDOWS 11

The optional update includes PC Game Pass

MICROSOFT'S FEBRUARY 2025 preview cumulative update (KB5052093) for Windows 11 24H2 comes with a range of bug fixes, enhancements, and features.

The update (OS Build 26100.3323) introduces a PC Game Pass referral card. This can be accessed by users with a connected Microsoft account by opening the Windows settings home page.

Microsoft originally introduced Game Pass for PC in 2019. Users pay a monthly subscription to access hundreds of games. Other perks include free trials of certain titles, as well as EA Play Membership.

There are also new functions for 'Narrator' scan mode. Users can press 'n' to automatically skip to the text after a link. The comma and period keys can also be used to jump to the start and end of items respectively. Multiple apps can now access the Camera application simultaneously to provide video streaming for both a sign language interpreter and an end audience.

Windows Spotlight has been overhauled to offer more information about background images. The Spotlight icon has also been tweaked and placed at the bottom right of the desktop. Users can now snooze or dismiss the 'Start Backup' reminder from File Explorer. Several bugs have also been fixed. URLs entered into the address bar now load correctly. The address bar also no longer overlaps files when users enter fullscreen mode (F11).

There are also fixes to some audio bugs. Previous versions of Windows would sometimes increase the volume to 100 percent when the device wakes from sleep. Another bug that sometimes caused USB audio devices to stop working after the PC was idle has also been corrected.

The KB5052093 update is currently optional, but the listed features will likely be rolled out more widely in the March 2025 Patch Tuesday release. **–ND**

The front tile design is open-source, so users can print their own.

FRAMEWORK REVEALS ITS MINI PC

Framework enters the desktop market

SINCE 2020, Framework Computer has been known for its range of modular, easily repairable laptops. The company has unveiled its first 'Framework Desktop'—a Mini-ITX size machine powered by AMD's new Ryzen Max series of mobile processors. This includes the AI Max+ 395. It boasts 16 cores, 32 threads, and a 40 CU RDNA 3.5 GPU.

The mainboard is a Mini-ITX form factor with ATX headers. It incorporates a PCIe x4 slot and set of rear I/O, including two USB 4 ports, two DisplayPorts, HDMI, and Ethernet ports. The system also supports a 256-bit memory bus, which theoretically enables up to 256GB/s memory bandwidth, though this comes at the price of upgradability. The LPDDR5x memory is soldered directly to the board.

The top-end Ryzen AI Max+ 39+5 configuration with 128GB of memory starts at $1,999. There's also a base Framework desktop that bundles an 8-core Ryzen AI Max 385 configuration with 32GB of memory, starting at $1,099.

Framework has developed a semi-custom 400W power supply with FSP for the desktop. It incorporates standard 120mm CPU fans with a thermal system. Buyers can also choose an RGB fan. The casing is also highly customizable, with a choice of black or translucent side panels, while the front panel can be made up of 21 colorful, customizable tiles. Users can even 3D-print their own if they wish.

Pre-orders are available now, with shipping due to start in Q3 2025. **–ND**

Twitch monetization tweaks

VIDEO-STREAMING SERVICE Twitch is on track to make some major changes to its monetization policies. According to CEO Dan Clancy, tools like subscriptions will be open to "most streamers", even if their account doesn't have 'Affiliate' or 'Partner' status. Currently, achieving 'Affiliate' status means meeting criteria like reaching 50 followers and maintaining an average of at least three concurrent viewers.

While newer users are excited about the potential for immediate earnings, established Affiliates and Partners have been more circumspect, feeling it could devalue the hard work they put into achieving their statuses. Streamers have also raised concerns that this could lead to ads on all Twitch channels. Currently, only monetized accounts like Affiliate channels display ads. Subscribers to Twitch's 'Turbo' service can also avoid pre-roll, mid-roll, and companion ads.

Twitch's latest move comes at a time when it has faced significant financial challenges. Despite Amazon paying $1 billion for the platform in 2014, it's lost around $2 billion in revenue every year since. In 2023, Twitch's chief customer officer, chief content officer, and chief revenue officer all left their positions, and the company laid off around 400 workers. In early 2024, it also laid off around 35 percent of its workforce. It remains to be seen if its new monetization policies will help the company to turn a profit. **–ND**

Google simplifies removing info

Since a 2014 European court ruling, Google has offered EU citizens the 'right to be forgotten'. This means they can have inaccurate or excessive information about them removed from search results.

However, this doesn't apply globally. Users have also complained that this request isn't easy to make. Google addressed this by overhauling its 'Results about you' feature. This provides more control for users to demand removal of data like names, phone numbers, and addresses. Users can click into search results to let them know what data is eligible for removal.

Google also offers an option to update outdated search results. If a search summary displays outdated information, Google can crawl the page to refresh the data. –ND

Gaming PC sales milestone

Research by IDC has shown that sales of handheld gaming PCs came to just under 6 million units in the last three years. The research firm also anticipates sales of 1,926,000 units for 2025, a figure that will likely be hit before the summer.

In 2023, sales peaked at 2.9 million units. This plummeted 50 percent in 2024 to just 1.5 million. Valve's Steam Deck is the most popular handheld, accounting for around half the sales in that time. IDC estimates that Valve has sold about 4 million units.

The data covered other devices like the Asus ROG Ally and MSI Claw, but didn't incorporate Chinese handhelds like the OneXPlayer X1.

These numbers pale in comparison to more mainstream devices. Sales of gaming laptops average 30 million annually. **–ND**

© FRAMEWORK, TWITCH BLOG

Jarred Walton

TECH TALK

Help us, AMD RDNA 4! You're our only hope

NVIDIA'S BLACKWELL ARCHITECTURE has landed with a bit of a wet splat. It's not that the GPUs are bad; it's that you simply can't find any of them in stock at reasonable prices. There's also the matter of the new cards offering relatively limited gen-on-gen performance improvements—unless you're specifically using Multi Frame Generation (MFG) as a smoothing substitute for higher frame rates.

AMD's answer is RDNA 4, which walks back from the GPU chiplet architecture of the prior generation. Specifically, AMD has taken steps to make its ray tracing hardware viable, and beefed up the AI compute potential. After the RTX 20-series launched in 2018, this is what we wanted from AMD. Better late than never, I guess.

To put things in perspective, the Navi 48 GPU at the heart of the RX 9070 XT has 64 RDNA 4 Compute Units (CUs) and a 256-bit GDDR6 memory interface with 16GB of VRAM clocked at 20 Gbps; nearly the same as the previous RX 7800 XT chip, with four extra CUs and a tiny bump to memory speeds.

On paper, that looks pretty mediocre: 6.7 percent more shader cores and 2.6 percent more bandwidth. But there's more here than the paper specs might suggest. For one, the boost clock gets goosed to 2.97GHz—a 22 percent increase over the 7800 XT that results in 30 percent more theoretical compute. But more important than the raw teraFLOPS are the tweaks under the hood.

RDNA 4's compute units offer substantial improvements to the shaders for up to 40 percent higher rasterization throughput. That comes via new out-of-order memory requests and dynamic register allocation. That's big, but not as big as the RT and AI gains. The ray tracing accelerators have doubled the ray/triangle and ray/box intersection rates, there's now dedicated BVH hardware as well as a new BVH8 compression scheme to reduce memory use, oriented bounding boxes, and better ray tracing hardware stack management. In practice, RT units should be about twice as fast.

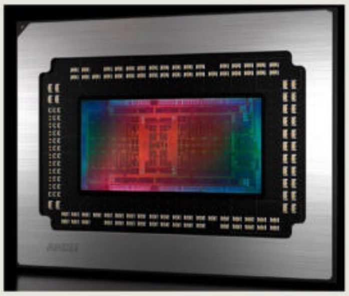

The Navi 48 marks a return to monolithic GPU designs for AMD, with 54 billion transistors.

© AMD

AI sees even bigger gains, as AMD doubled the FP16 throughput and supports sparsity for up to 4X the compute. AMD also added native FP8, INT8, and INT4 support to the AI accelerators. FP8 doubles throughput relative to FP16 without compromising much on precision. With FP8 and sparse matrices, RDNA offers 8X the AI compute of RDNA 3.

In practice, these combine to allow the RX 9070 XT to deliver better ray tracing and AI performance. Stable Diffusion XL performance, for example, was 39 percent faster on the 9070 XT. The 9070 XT offered 10–12 percent higher ray tracing performance across 1080p, 1440p, and 4K resolutions. Rasterization performance was only 3–5 percent worse than the previous halo card.

AMD is aggressively pricing its RDNA 4 GPUs, with a $599 MSRP for the RX 9070 XT that only ends up 3–5 percent below the new RTX 5070 Ti. The RX 9070 in turn has a $549 MSRP, ostensibly the same as the RTX 5070, with 33 percent more VRAM capacity and 4–8 percent higher performance.

None of that matters if you can't buy the cards. Current and previous generation GPUs are basically out of stock for all but the budget RTX 4060, RX 7600, and Arc B570 cards. Cards from third-party resellers are typically 50 percent or more above MSRP. AMD's new RDNA 4 GPUs sold out at launch within minutes.

We'll need to wait and see, but it feels a lot like 2021, where cryptocurrency-induced GPU shortages coupled with pandemic PC upgrades resulted in horrible retail GPU prices. This time, it's the demand for AI causing shortages, as Nvidia has seemingly abandoned gamers in favor of the data center market—88 percent of its record $130 billion in revenue last year came from data center and AI, compared to just under nine percent from gaming.

Hopefully, AMD will keep the supply of RDNA 4 GPUs coming.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

THE LIST

THE BEST MINI PCs

In recent years, mini PCs have bulked up to become modest yet genuinely viable gaming platforms. Sure, if your chief requirements are top frame rates and detail settings beyond 1440p for the lowest outlay, you are better off self-building. But if you're after a tiny, off-the-shelf PC with awesome everyday computing power and 1080p to 1440p gaming capabilities—well, you've come to the right place.

5 BEELINK SER9

It might be small, but inside this all-aluminum machine is the most powerful iGPU we've tested to date: the AMD Radeon 890M. It runs quietly and looks great, no matter where you put it. It's an impressive performer, but you'll have to pay to the tune of about $1,000. Ouch. It also has a beautifully svelte design with a chassis that's stacked with ports, including USB Type-C, Type-A, HDMI, DisplayPort, and Ethernet. **$999 www.bee-link.com**

4 AYANEO RETRO MINI AM02

With its NES-aping design, the Retro Mini AM02 is the cutest mini PC around. But that beauty is not just skin-deep, because inside lurks the excellent AMD Ryzen 7 7840HS APU and a level of performance customization few other machines offer. It can definitely get a little 'chatty' when its fans crank up, should you start making that AMD CPU run harder. That aside, there's plenty to like, including good I/O and a dinky touchscreen. **$699 www.ayaneo.com**

3 GEEKOM AX8 PRO

It's stunning how much performance you can pull from a PC so small. The Geekom AX8 Pro is incredibly dinky, yet the AMD Ryzen 8945HS APU with a Radeon 780M iGPU at its heart delivers some excellent iGPU-based gaming performance without sounding like a jet turbine. This isn't the most powerful mini PC, but if you want the smallest, most svelte square, it's an excellent choice. **$899 www.geekompc.com**

2 ASUS ROG NUC 970

If you're after a mini PC with the best gaming performance, this is the one to covet, but you'll need to be prepared to part with a significant amount of cash. Along with the Intel Core Ultra 9 185H CPU, you get NVIDIA's RTX 4070 Mobile GPU. It's not up there with the desktop 4070 for raw performance, but if you want a performant mini gaming PC that delivers at 1440p, the ROG NUC is for you. **$1,629 www.asus.com**

1 MINSFORUM ATOMMAN G7 PT

If the best mini PC needs to be able to game, thanks to the RX 7600M XT GPU, the AtomMan G7 PT is capable of delivering. Its eight-core, 16-thread Zen 4 chip is equally capable of productivity goodness. The CPU and GPU are soldered in, but the AtomMan G7 PT is a great machine. Just bear in mind that you'll need your own DDR5 RAM, M.2 storage drive, and Windows key. **$999 www.minisforum.com**

© BEE-LINK, AYANEO, GEEKOM.

Jeremy Laird

TRADE CHAT

Stuck between a rock and a tariff place

THE CEOS OF NO LESSER CONSUMER TITANS than Best Buy and Target both think that tariffs are going to turn into higher prices for consumers. That's right—just when you thought graphics card prices couldn't get any worse, here comes something to mess things up all over again.

Price increases for American consumers are likely. So said Best Buy CEO Corie Barry in response to the recent reactivation of 25 percent tariffs on goods and services from Canada and Mexico, plus an extra 10 percent on China for good measure, adding to the 10 percent tariff the Trump administration applied back in January.

Barry said China and Mexico are Best Buy's top two import territories, with 55 percent and 20 percent of its products sourced from those countries. Meanwhile, the CEO of Target, Brian Cornell, said tariffs mean "prices will go up".

Cornell was speaking specifically about the impact of tariffs on Mexican goods. Admittedly, Target isn't exactly a PC component specialist, selling only fully built PCs and laptops, so Cornell's comments aren't as directly relevant as Best Buy's observations. But Cornell is nothing if not wired into the impact of tariffs on consumer prices.

"We're going to make sure that we do everything we can to protect pricing. But if there's 25 percent tariff, those prices will go up," Cornell said. If the overall direction is clear enough, the details are arguably less so. "We've never seen this kind of breadth of tariffs, and this of course impacts the whole industry. So it's not just a Best Buy question, it is a broad industry question. And I say that because that makes the estimation of the impact all the harder," Best Buy CEO Barry explained.

The impact of this on graphics cards in particular will be hard to gauge. GPU pricing has been volatile for years now, so if there's another spike in prices, it will be tricky to say that it's definitely a direct consequence of tariffs. In the short term, a load of new Nvidia RTX 5070 and AMD RX 9070 and 9070 XT graphics cards will have been shipped to the US for early March, comfortably before the 10 percent China tariff activated on March 4.

> It's certainly possible that some or even all of the tariffs may be revoked

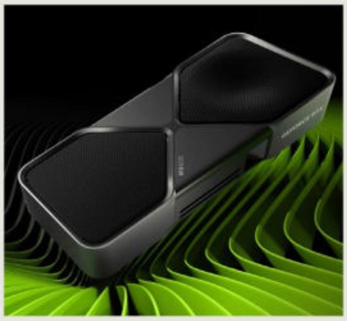

If you spot one of these anywhere near MSRP, jump on it!

© NVIDIA

The prices of those GPUs won't be directly impacted by the tariffs. But what happens next? It's certainly possible that some or even all of the tariffs may be revoked. Of course, importers and wholesalers know that, so might they hold fire on bringing more cards in for a while? If so, that will mean tight supply, which will push prices up. You can see the problem.

The bottom line is that none of this is good news for sensible prices. Let's say a whole mountain of graphics cards are imported from China in the near term to boost supply. Well then, they'll be hit with 20 percent tariffs, and the bulk of that will no doubt be passed on to gamers. On the other hand, if importers try to wait out the tariffs, then supply will be constrained. Guess what? That pushes prices up, too.

All of which means everything is pointing in one direction, none of it good. If you're even remotely thinking about snagging a circa $500 to $600 graphics card, and you spot Nvidia's RTX 5070 or AMD's RX 9070 or 9070 XT anywhere near their purported MSRPs, jump on it. I say that even about the Nvidia RTX 5070, which is a disappointing card, but unlikely to get any cheaper any time soon.

Right now, then, the market is stuck between supply issues and tariffs. Fix one, and you get hit by the other. It's all more than a little disappointing and frustrating. Just as AMD's new graphics cards threaten to make the market more competitive once again, and even Nvidia has been talking up RTX 5070 supply, the blunt impact of tariffs could drive up prices all over again. What a bummer.

Six raw 4K panels for breakfast, laced with extract of x86... Jeremy Laird eats and breathes PC technology.

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Cloning M.2 drives
> Rename files quickly
> Fix sync errors

Upgrade NVMe drive

I want to swap out my M.2 NVMe 256GB boot drive for a 1TB one, as I'm constantly running out of space. I have one eye on a future PC upgrade, so I'm looking for a drive that will move with me. Can you recommend any models, and how would I go about transferring my Windows install from one to the other? **—Bryce A Lawson**

THE DOCTOR RESPONDS: Bryce's mobo only supports PCIe 3.0, which maxes out his read/write speeds at around 3 Gb/s. However, most PCIe 4.0 drives, which can read/write at over 7.5 Gb/s, are backward-compatible with PCIe 3.0. There is a premium for those top speeds—Samsung's 990 PRO 1TB costs around $116 on www.newegg.com. If that's too rich—particularly as you won't get better than 3 Gb/s until you upgrade—then a range of PCIe 4.0 1TB drives that top out around 5 Gb/s costs around $65, including Crucial's P3 Plus, or you can choose a budget-friendly option via Silicon Power's UD90 range—its 1TB drive can be had for just $52.99 on newegg.

The question is how you clone the old drive to new—

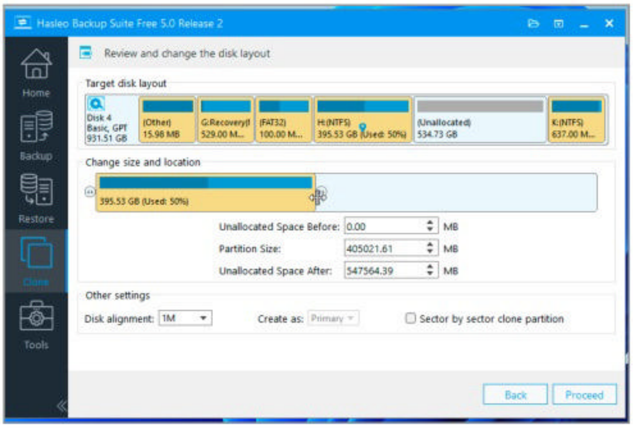

Unallocated Space Before: 0.00 MB
Partition Size: 405021.61 MB
Unallocated Space After: 547564.39 MB
Disk alignment: 1M
Create as: Primary
☐ Sector by sector clone partition

Use Hasleo's clone tool when upgrading your boot drive.

© HASLEO

check your mobo's specs to see if it has a spare NVMe slot. If so, fit the new drive, then use a free drive-cloning tool like Hasleo Backup Suite (www.easyuefi.com) to clone the old drive to new while booted into Windows. To do this, select Disk Clone under Clone. Select your old NVMe drive from the list, click Next, then select your new drive (identifiable by its size and being listed as 'unallocated'), and click Next again.

After that, you'll be given the opportunity to review and change the disk layout—in other words, expand the size of your drives to make use of the new capacity. Start by selecting any drives to the right of your main drive (ignore the fact that it'll be allocated a different drive letter), then drag them to the right of the 'Change size and location' box. This should place the unallocated space next to your main drive—ignore the fact that it's being allocated a different drive letter. Windows will resolve that issue when you shut down your PC, and remove the old drive before rebooting.

Now, select your main drive, and expand its size using the drag handle—you can give it all the available space, or reserve some for another partition. When you're happy, click Proceed. There's no confirmation dialog—Hasleo will begin the cloning process. It should be a swift operation—no more than a matter of minutes—after which you should shut down your PC. Remove the old drive and, if necessary, place your new drive in its slot (assuming, like the Doc's mobo, one NVMe slot is faster than the other) before rebooting.

Be prepared to wait on that first boot while your PC adjusts to the new setup—you may have to wait up to 10 minutes before everything settles down. Once done, you should find your new drive up and running. You'll no longer have to worry about running out of drive space—for now, at least.

Batch renaming tools

I've got a lot of files I need to rename, and I've been looking at various batch tools. I don't want to spend any money if I can avoid it—can you recommend one?

—David B Tucker

THE DOCTOR RESPONDS: The Doc's go-to batch renaming tool is Bulk Rename Utility (www.bulkrenameutility.co.uk), which is free for personal use. It's reasonably simple to use, but offers

↘ submit your questions to: doctor@maximumpc.com a wide range of rename options. These include the standard Replace, Remove, and Add options, along with more sophisticated tools like Move/Copy Parts, Append Folder Name, and adding dates or numbers. There are also filters and a RegEx tool.

What sets Bulk Rename Utility apart is the fact that you can combine these options to create really sophisticated rename operations. Making things simpler is the Preview window where you not only select which files will be renamed, but also see how your naming setup will affect them before you click Rename. Crucially, you can craft multiple rename patterns, and save them as Favorites for quick recall.

The feature isn't well advertised, but Bulk Rename Utility can also rename files using a CSV list, whereby you specify the source filename and the name you'd like to change it to. This is perhaps one of the program's weaker features—when the Doc wants to rename a group of files using a list, he prefers Advanced Renamer (www.advancedrenamer.com). This lets you create and edit lists within the program itself, then match them up using drag and drop. It's also free for non-commercial use.

Syncthing error

I'm using Syncthing to keep large folders in sync between my server and QNAP backup NAS. Recently, I upgraded one of the NAS drives, so had to remove and reshare the folder link from my server. I attempted to set up the new drive with the same folder structure as before, but my first problem came when I received a 'folder marker missing' error after recreating the folder link on the NAS. I've had one of these errors before, so knew I needed to create a new hidden .stfolder folder in the directory on the NAS. However, after leaving it to

© BULK RENAME UTILITY

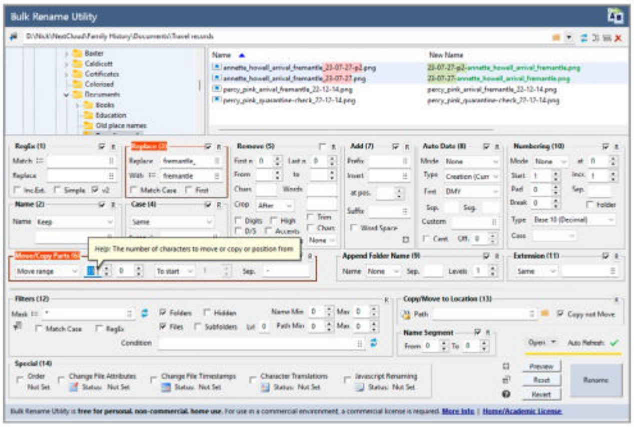

Rename files quickly with the help of Bulk Rename Utility.

copy the files back to the NAS, it had thrown up a new 'Out of sync' error claiming the global state was twice the size of the local state. I checked my server, and it showed both global and local state were the same.

I've tried removing and recreating the folder link, but it just produces another .stfolder error, which I resolve, only to find the reported global state gets bigger each time. How do I fix this problem? **—Lou Britt**

THE DOCTOR RESPONDS: This appears to be the combination of two separate errors leading to major confusion on Syncthing's part. Each time you manually recreate the hidden .stfolder folder to fix the error with the missing folder marker, Syncthing views all the files on your NAS as different to those on the parent folder, and adds them to the ever-growing list of out-of-sync items, hence a 1.79TB folder quickly claims to be three and then four times the size.

Fortunately, you can fix this without deleting the entire folder. First, open Syncthing on your NAS, and remove the shared folder again—this removes the .stfolder folder while leaving the previously synced files intact. Now, re-share the folder again—either by restarting Syncthing on your NAS to trigger the prompt to add it back in, or by unsharing it and then resharing it from the instance of Syncthing running on your server.

When prompted to set up the share again, follow the usual procedure, but make one crucial change: select Receive Only under Folder Type on the Advanced tab before clicking Save. The first difference you'll notice is that the .stfolder marker is created automatically, so you avoid the Stopped error message. The second thing is that you'll immediately see the Out of Sync Items figure reappear and climb as the files are scanned. This time, it's accompanied by a 'Revert Local Changes' button. Wait for the scan to complete, then click the button. Syncthing should sort itself out, and report that the folder is up to date. You can then go into Edit > Advanced tab and change Folder Type back to Send & Receive, if you so wish.

Disable Protected View

Whenever I open an Office document downloaded from the internet, I get the Protected View warning. I could live with this, but lately, Word seems to take an age to process the file before displaying the message. What's the easiest way to bypass this? **—Fred E James**

THE DOCTOR RESPONDS: We'll assume you're obtaining the files from a trusted source—say, a cloud drive—and have no concerns over the content. If in doubt, a quick scan with your antivirus before opening the file should determine whether it's safe. You can now disable Protected View, which needs to be done individually for each Office application you use: Word, Excel, PowerPoint, and so on.

Open each application in turn and, if necessary, switch to the File tab. Click Options, select Trust Center in the left-hand pane, and click Trust Center Settings to open another window. Select Protected View in the left-hand pane, and you should see three 'Enable Protected View...' options. In most cases, disabling the first ('Enable Protected View for files originating from the Internet') will be sufficient—you may also want to disable Protected View for Outlook attachments if you use Outlook, but we recommend leaving 'Enable Protected View for files located in potentially unsafe locations'. Click OK twice to enable the change.

Container won't stop

I'm trying to shut down a Podman container so I can update the volume in Debian, but when I try to stop it via Cockpit or the Terminal, it keeps respawning. Please help! **—Federico Kelly**

THE DOCTOR RESPONDS: This will be linked to two things: a 'restart unless stopped' instruction in the Podman run command, and the service setup info used to launch the container at startup. In Cockpit, navigate to Services and switch to User view, where you should see a list of the containers you've configured. Click on the troublesome container to access its status page, then flick the switch off. This will stop the container, allowing you to update the image. Once done, flick the switch back on to restart it.

TINY BUT MIGHTY

You still can't buy that GPU, but this ITX Adonis delivers some astonishing 4K performance

LAST MONTH, we had a fantastic reader question surrounding our lack of AMD ITX builds. We've done numerous Intel ones, but Team Red was nowhere to be found. The reason for that: a lack of motherboards. AMD systems don't naturally lend themselves to ITX investment. In fact, you can count the ITX boards for the AM5 platform without running out of fingers (at least the ones we can request in).

Still, it got us thinking about how much power you could slip into an ITX rig these days, particularly with 9th gen Ryzen and Nvidia's latest 50 series. With the launch of Team Green's latest twin-slot Founders Edition cards, it really got us thinking. We grabbed our Ryzen 9 9900X, 32GB of 7400 MT/s memory, one of Nvidia's latest RTX 5080s, and a Fractal Design Era 2, and started pulling the parts together to build something truly potent.

Did we succeed? Yep; this thing's an absolute goliath. From insanely high frame rates, to dominant CPU performance and astoundingly respectable power and temperatures, it just dominates in every arena, with a remarkably small footprint.

Unfortunately, like with every other Nvidia GPU right now, you just can't get hold of the things. Nvidia is apparently working towards correcting that supply issue, but we're talking months out, as the company remains fixated on its main money maker: AI server-grade componentry. By then, tariffs will likely have had quite a dramatic effect on pricing. Still, it's hard to tell how exactly this will all play out until they land back with us in stock. Until then, welcome, and we hope you enjoy our ridiculously tiny ITX monster machine, built in the lovely Era 2 chassis from that Swedish design house we all know and love. **–ZAK STOREY**

BENCHMARKS

Part		Price
CPU	AMD Ryzen 9 9900X	$382
Motherboard	Asus ROG Strix X870-I Gaming WiFi	$518
CPU Cooler	NZXT Kraken Elite 240 RGB	$254
RAM	32GB (2x16GB) Corsair Dominator Titanium DDR5 @ 7400 C34	$190
SSD	2TB Corsair MP700 Elite M.2 PCIe 5.0 SSD	$190
GPU	Nvidia GeForce RTX 5080 Founders Edition 16GB	$1,000*
Case	Fractal Design Era 2 ITX	$200
PSU	850W Phanteks Revolt SFX 80+ Platinum	$189
Total Price		**$2,923**
Core Price*		$2,280

*"Core Price" refers to the key components generating performance (CPU, GPU Mobo, OS SSD, RAM), not accessories. GPU price based on RRP at time of initial launch.
*RRP listed, currently retail prices range from $1,680 - $2,200. Likely until stock returns.

PRICES CORRECT AT THE TIME OF PUBLISHING

https://cdn.jwplayer.com/videos/QPWcX01N-u2lN49He.mp4
Please type this URL into your browser if the link is broken

STEP-BY-STEP GUIDE
PG.22

NZXT
GEFORCE RTX

AMD
RYZEN
AMD Ryzen 9 9900X

CPU

AMD Ryzen 9 9900X

It has long been a staple part of our modern-day testbeds. AMD's Ryzen 9 9900X toes the line between too power hungry, toasty, and not having enough multi-threaded grunt. Based off AMD's Zen 5 architecture, on the AM5 socket, AMD's packed this thing with 12 cores and 24 threads clocking along nicely at 5.6 GHz at max boost. All of that built out on TSMC's 4nm FinFET process tech.

It's that pricing that really delivers for us, though, and keeping it that low right now is a blessing in disguise. Yes, there's some argument if gaming is your only focus that the Ryzen 7 9800X3D might have made more sense, particularly given the X3D cache, but we've found it to be often inconsistent with the titles it supports. Given the slightly lower clock speeds, and fewer cores, we'd rather have the consistency/jack of all trades philosophy that the 9900X provides instead.

$382, www.amd.com

STEP-BY-STEP GUIDE
PG.22

Motherboard

Asus ROG Strix X870-I Gaming WiFi

One of the reasons we've stayed away from AMD ITX builds is the pricing on those motherboards, and Asus is showcasing exactly that with its X870-I motherboard here, coming in at an eye-watering $518. That's some serious outlay, even for a board as well-kitted out as this.

Asus has packed this thing with a 10+2+1 power stage design, support for DDR5 up to 8600 MT/s (more on that in a moment), and twin M.2 slots (one at 5.0 and one at 4.0). Connectivity is fairly average too, there's six USB Type A ports on the back at varying speeds, along with three USB Type C, a 2.5Gb Ethernet port, plus WiFi 7 and an HDMI out, and that's about it.

$518, www.asus.com

FOR THOSE
ROG STRIX
WHO DARE
ROG
STRIX
STRIX X870-I GAMING WIFI

RAM

32GB (2x16GB) Corsair Dominator Titanium DDR5 @ 7400 C34

Hold up—is that stock standard Intel XMP memory we've got there? Yep! That's right, unlike many of the DDR5 kits on the market, the Titaniums have specific kits for both AMD and Intel.

DOMINATOR
TITANIUM
DDR5
DOMINATOR
TITANIUM
DDR5

But did you know that you can still run Intel XMP memory on AMD-specific motherboards? We chatted long into the night with Corsair's memory nerds about whether this would work with our build or not, and the answer is simply 'yes'.

With that, we've got an awesomely quick 32GB kit of 7400 MT/s Dominator Titanium. It's a little overkill, admittedly, particularly as you can't technically see it with the Era 2's surrounds encapsulating it, but that sub 10-second real world latency performance is just too juicy to pass up.

$190, www.corsair.com

CPU Cooler

NZXT Kraken Elite 240 RGB

NZXT's latest coolers are stunning to look at, and we're serious fans of what it has done with the design. Our Fractal Era 2 has capacity for a 240mm AIO in the roof, and that's what we've gone with here—a perfect pairing for that Ryzen 9 9900X.

With a massive 640x640 2.72-inch IPS display, beautifully braided tubing, and a cable-less design, it's a joy to use and super-simple to install. NZXT has integrated the cable connectors directly into the radiator, meaning no more awkwardly routing cables around VRM heatsinks from the AIO. Better still, it's powered by NZXT's own custom-built and designed 'Turbine' pump solution, producing some seriously quiet cooling, and an interesting step away (or should we say around) Asetek's patent dominance in this arena.

$354, www.nzxt.com

SSD

2TB Corsair MP700 Elite M.2 PCIe 5.0 SSD

Another Corsair product, this time hailing from its NAND Flash division. To keep things simple, we've opted to go for a single drive on our ITX power house, this time with Corsair's latest MP700 Elite M.2 PCIe 5.0 SSD. It's not by any means the fastest drive out there, but it has two advantages we've chosen it for.

Firstly, it has got a removable heatsink (secured down with only Phillips screws), and secondly, it's stupidly efficient, thanks to Phison's latest E321T controller. Sequentials sit at around the 10GB/s mark on the read, but owing to its DRAMless design and that controller, typical 85 C temps have dropped by 30 degrees or more, which is ideal for a compact system like this.

$190, www.amd.com

STEP-BY-STEP GUIDE PG.22

GPU

Nvidia GeForce RTX 5080 Founders Edition 16GB

What to say about Nvidia's latest 5080 card (other than you can't buy it)? Well like its 5090 sibling, it touts that twin-slot cooler design (admittedly rather tame in appearance). That's perfect for our Era build, which will really only support cards of that thickness.

On top of that, Nvidia has crammed this with 10,752 CUDA cores, 336 TMUs, 84 ray tracing cores, and 336 Tensor cores. We haven't reported on the ROPs as well; there's a lot of drama surrounding that, mostly around false reporting, missing ROPs, and more. Gamers Nexus has done a fantastic video on it.

Still, our 5080 sample is a perfect fit for our Era setup, and will deliver comfortable performance at 4K. One thing to bear in mind: it's still on a 5nm process (slightly more efficient than the 4080 Super), and has only seen a five percent hardware increase on last gen, so outside of memory improvements, don't expect massive upticks in performance if you're coming from a 4080 or similar.

$1,000, www.nvidia.com

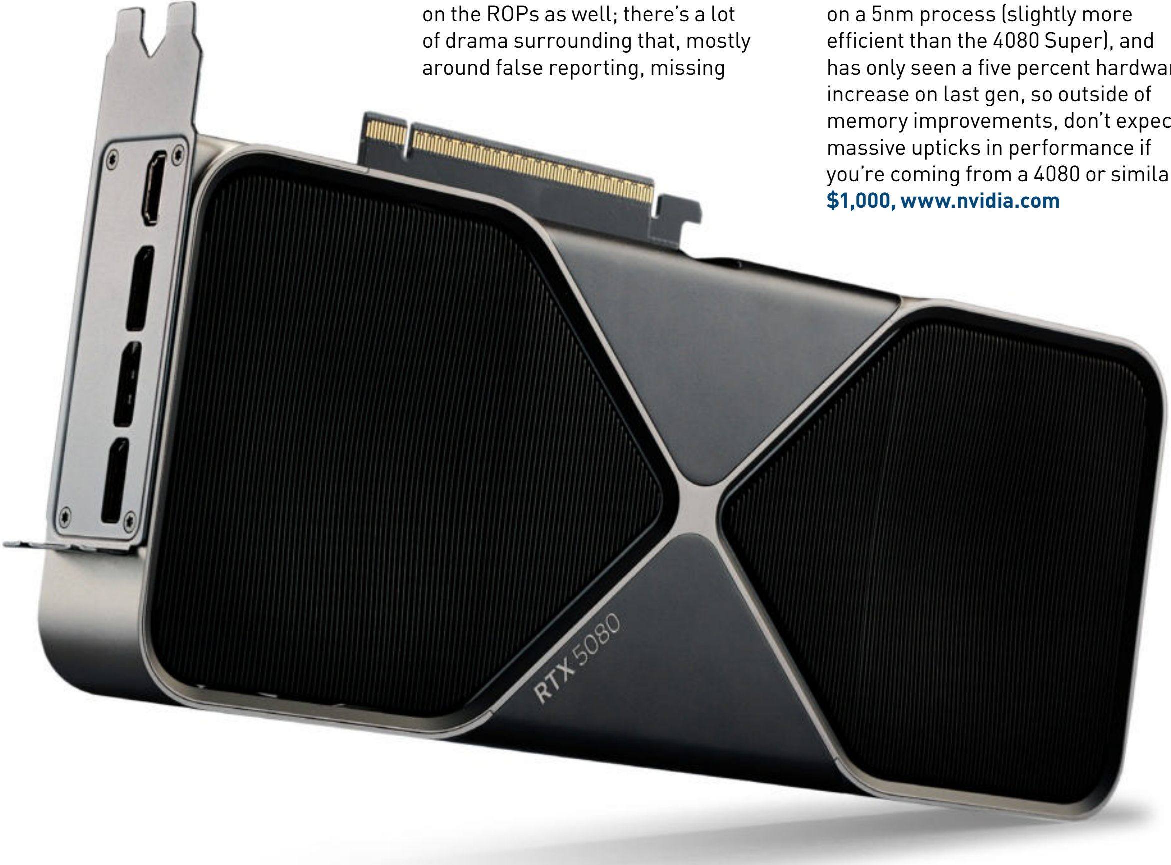

PHANTEKS
REVOLT SFX
PHANTEKS
PHANTEKS REVOLT SFX

PSU

850W Phanteks Revolt SFX 80+ Platinum

Following on from our chassis decisions above, and with plenty of experience with the Terra for our power supply we knew we wanted to stick with an SFX unit, and we knew it had to be on the ATX 3.0 platform with a Gen5 12VHPWR cable built in as standard. Oh, and have plenty of juice for future upgrades. Surprisingly, there's not a whole host of power supplies out there that'll fit the bill, so we grabbed this Revolt SFX from Phanteks, which, despite its size, and rating is surprisingly affordable. Better than that, though, all of the cables are short, flat ribbons, making builds like this even easier.
$189, www.phanteks.com

Case

Fractal Design Era 2 ITX

The Era 2 is a remarkably beautiful chassis that needs to be seen in the flesh to be fully appreciated. The design language, curves, wooden accents, perforations, paint job—it's an outstanding little ITX chassis, not too dissimilar to the Terra we built in late last year. The major difference? Other than some internal design shifts here and there (it still features a shifting motherboard tray, giving you adaptable clearance for GPUs and CPU coolers), there's a ton more roof space, meaning that you can fit radiators up to 280mm in there.

That said, there are a few caveats. Make sure you get an SFX PSU, not an SFX-L, and we recommend sticking with a 240mm AIO, not the full-fat 280mm (Fractal agrees with us on that one too). Available in black, white, or blue satin-coated finishes, it's a true treat, but requires an expert builder to get to grips with.
$200, www.fractal-design.com

THE UNDERDOG

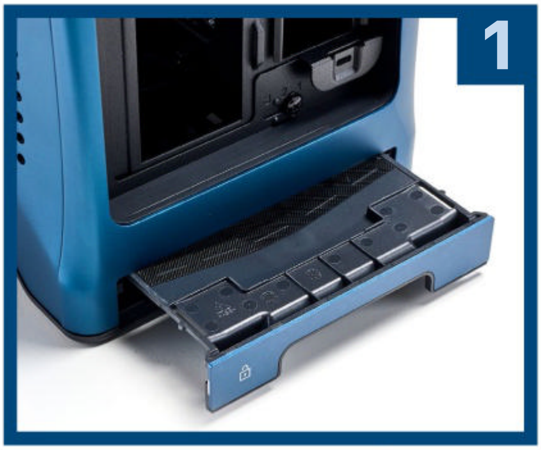
1

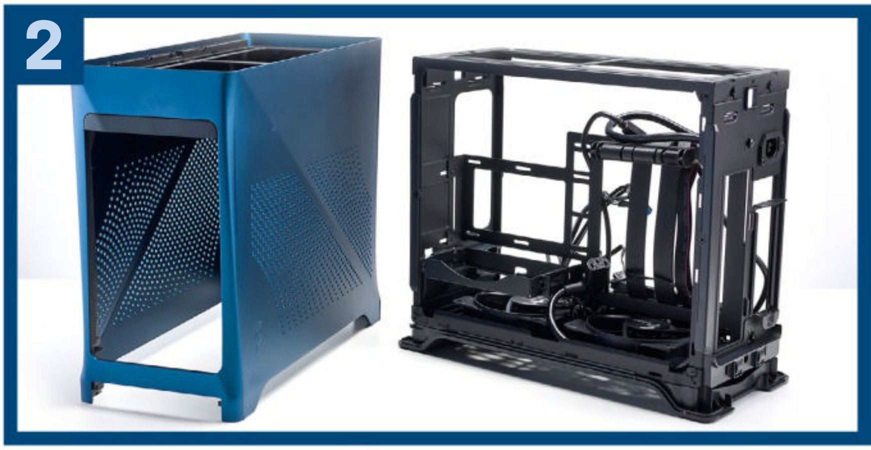
2

3

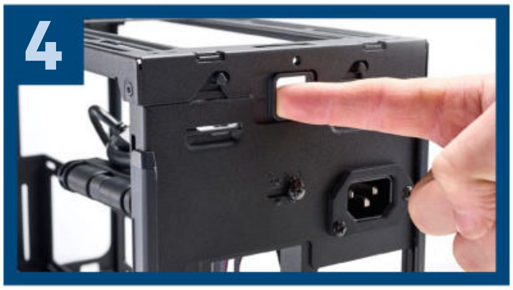
4

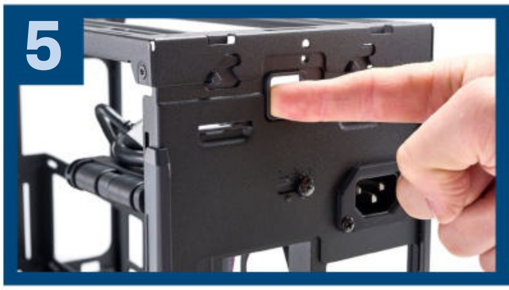
5

LENGTH OF TIME: **2-3 HOURS**
LEVEL OF DIFFICULTY: **HARD**

WE LOVE A challenging build, and nothing comes close to a compact ITX system like this. There's no denying this one, not one bit. If you're building in the Era 2, you're building a true SFF build. No fluff, no extra space inside; every cable route needs to be planned carefully, the components perfectly selected, and the build order needs to shift back and forth. Adapt, overcome, and improvise. The Era 2 is exactly what you expect it to be: super-slick, and jam-packed with features to make even the expert system builders out there happy, but it demands the excellence of its builder. It'll be frustrating at times, and challenging at others, but the satisfaction you get at the end of it, and the system you're left with—well, it's an unparalleled experience.

We've built in something similar from Fractal before, namely the Terra ITX chassis. It's actually slightly older and smaller than this case. The two of them, however, are built around an incredibly similar internal layout. You've got a shifting frame that allows you to adjust the depth between the two sides. If you've got a thicker graphics card, you can shift the internal frame over and give yourself more room for it while limiting the height of your CPU cooling solution instead. In the case of the Era 2, that's actually somewhat negated, as the 240/280mm radiator support in the roof allows you to fit an AIO, which, outside of Asetek's new pump/fan AIO units, are generally quite low profile by design.

Otherwise, internal footprint and desk footprint are remarkably similar, bar the extra height. One major difference comes in the actual frame itself. Unlike the Terra with its gullwing panels, the external frame of the Era 2 is actually a shroud; an enclosure that sits around the overall skeleton, locked in place by an incredibly robust dust filter mechanism found at the rear of the chassis.

We'll sign off this segment by saying simply that this is a challenging build—perhaps one of the most difficult we've done in the last year or more. That's not a slight on the Era 2; it's the constraints you're faced with building inside a true SFF chassis, which you need to take into consideration.

DIVING IN THE DEEP-END

Of course, as with all *Maximum PC* builds, we must first start with the chassis teardown. Head around to the rear of the case, and you'll spot a small, rectangular piece of plastic, powder-coated in the same finish as the rest of the chassis with a padlock symbol on it. Pull this out all the way, and set it to the side. **[STEP 1]** This will effectively 'unlock' the internal shell, giving us access to the skeleton within.

With the dust filter removed, we're also now going to remove the topmost wooden filter (actually removable, would you believe it?). It's held in place via the power of gravity, magnetism, and a bevelled edge) push down on one end of it until it tilts upward. Then you can grab this and take it out, placing it to the side.

With that, you can then carefully lift the exterior shell off the Era 2's core chassis, revealing the skeleton underneath **[STEP 2]**. You can see from our snazzy little shot there: the two segments side-by-side. We'll have already built in this once by the time you see this shot (for a separate article), so the cables are a little messy. To give you the low down, we've got that PSU hanger in the front, with a power cable passthrough running to it from the rear of the chassis. That runs around the motherboard tray. The PSU 'holder' slides out on a rail mechanism easily enough. This is one of the few oversights on the case, as it locks into place on a latch, which is great, but it does mean that the PSU is fixed to the side of the case, and there's no adjusting the depth of it without some 'bodging'. That'd be handy if, say, you're using something like be quiet!'s

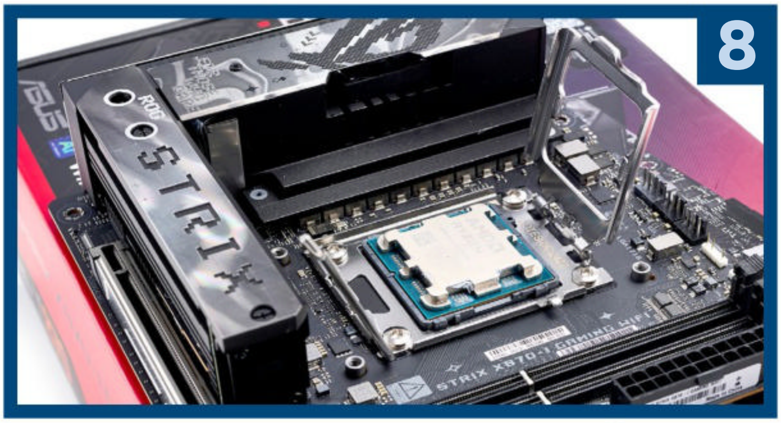
8

9

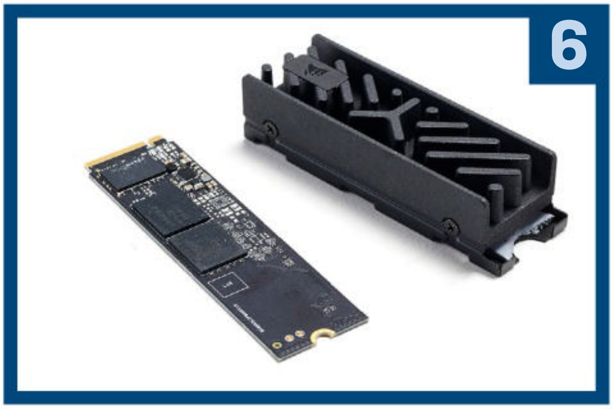
6

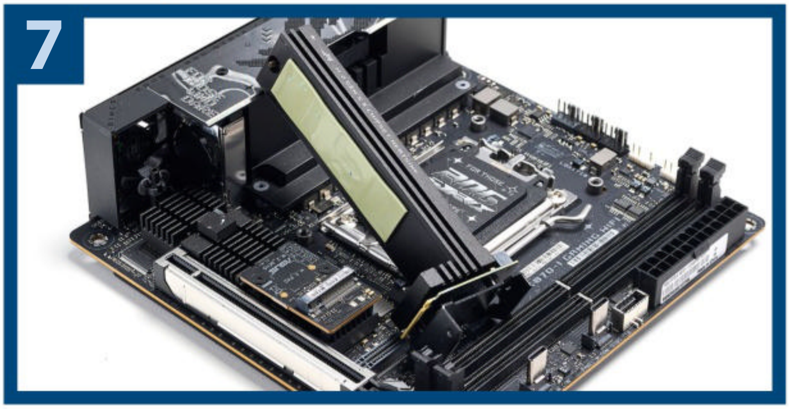
7

SFX line, as the orientation of the power port means the passthrough cable bulges out against the side of the case. **[STEP 3]**.

You can also see a PCIe 4.0 riser cable there to allow us to install the graphics card on the other side of the chassis, and Fractal's also included twin 120mm fans in the base of the thing, providing us with plenty of elevated cool intake air directly up into both compartments, as well as cooling the PSU. It's all cleverly planned out, taking into consideration the limited thermal capacity the overall system has.

Speaking of thermals, that top radiator mounting location is thankfully a removable bracket with one of the most ingenious locking solutions we've seen. Fractal has effectively implemented a latch system that you can access via these two finger tabs on either side of the case, front and back **[STEP 4/5]**. Slide the latch across, and the two tabs will move over, and you can pop the radiator bracket up and off. To secure it down, you can simply push it down and lock it into position. Nice.

You can also see that the skeleton is going to cause a bit of drama. Similar to our Terra build, you need to think about your cable management way ahead of time, particularly CPU power and front I/O. This is a case that really benefits from running cables early, before any components are in. Fortunately, unlike our Terra build, we're using a much smaller (yet way more potent) SFX PSU, which should give us a touch more clearance to play with.

THE HARD WORK BEGINS

Now our chassis showcase is out of the way, as always, we're going to move on to preparing the motherboard. First up on the chopping block is removing the heatsink from our Corsair MP700 Elite. To do that, take your finest small Philips head and remove the four screws on either side of the heatsink. Pop them apart, and there's one whole M.2 SSD stick, ready to go for your pleasure.

While we have it apart, we'll give you a little lesson on the drive composition. The chip closest to the connectors at the very end there is the controller; the two chips further up the stick are actually the NAND packages that house all the data. As you can see, there are no additional smaller chips, which signifies that there's no DRAM cache or similar. It's handy to know if you're not sure what your SSD setup actually is. Typically, drives with DRAM cache (that aren't QLC) usually provide more consistent performance over time, while a drive with a DRAM cache will run quickly at first until that cache is full, then performance will slow dramatically **[STEP 6]**.

Moving on to the motherboard, it's time to do the world's most ridiculous masonic handshake to get to the M.2 PCIe 5.0 slot. It's situated closest to the motherboard underneath a rather beefy heatsink. The heatsink itself effectively comes in two parts, and has an interconnect that runs up it. The TL;DR is effectively that the top-most M.2 port is the PCIe 4.0 one (underneath the first heatsink), and the bottom-most M.2 port is the PCIe 5.0 port. To get to the second, much faster port, you need to remove two tiny Philips screws attached to the bracket on the bottom of the entire unit, only accessible once you've removed the top heatsink. Ridiculous doesn't quite cover it, but you do get a chunky heatsink to help eliminate any excess thermal buildup. It's just a nightmare to put back together **[STEP 7]**.

As always, line the drive up with the M.2 slot, align the notches, and click it into place. There's no securing mechanism here—just the power of the heatsink being secured back down. Do what you did in reverse here, and make sure you remove the thin plastic film from the thermal pad, before securing it back down on top of your 5.0 drive, and you'll be set.

With the SSD in position, we can now move on to our CPU installation. As usual, we lift the retention arm out and up at the side, then lift the bracket up

10

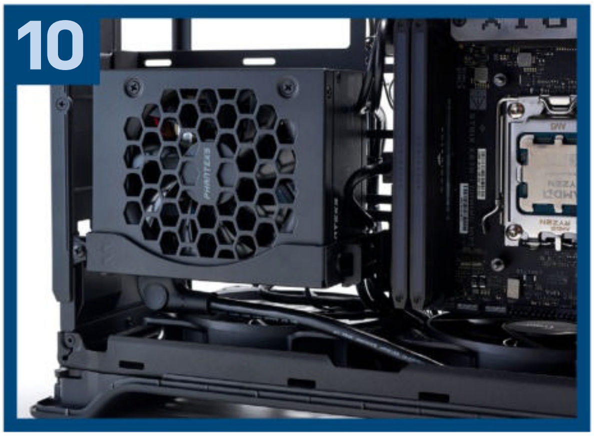

11

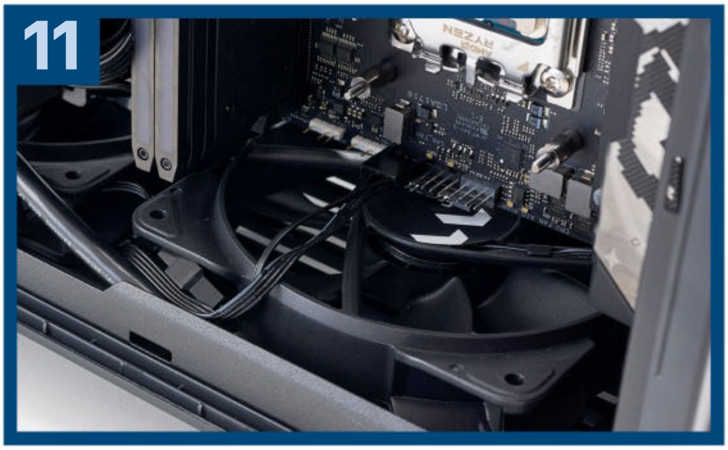

12

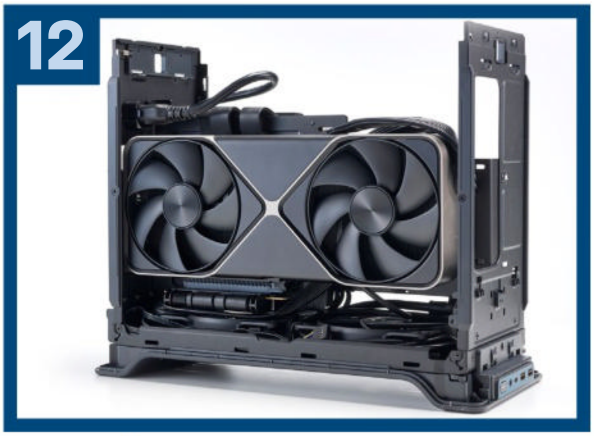

13

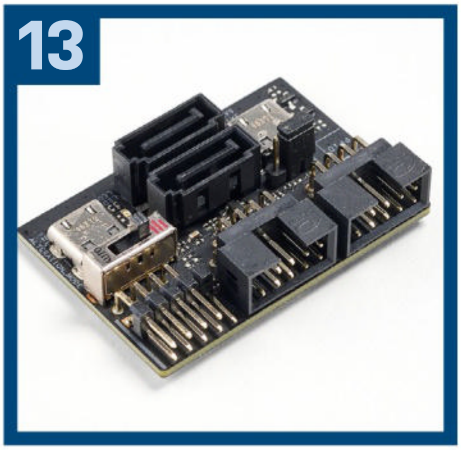

14

15

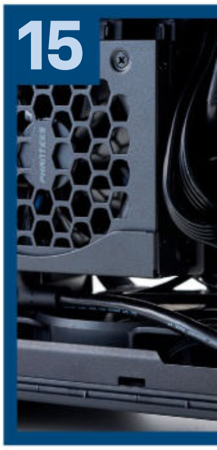

before carefully placing the CPU into position, making sure to align it with the notches, and the writing, and the golden triangle located on the chip, the socket, and the bracket. Place it carefully into position, lower the bracket back down, and resecure. We've proactively removed the plastic cover here, but you don't need to as it'll pop off as you secure the bracket back down **[STEP 8]**. We've also removed the cooler mounting brackets off the backplate to allow for our next step too. One thing you might also be able to spot are those two oddly shaped USB Type C ports sorta sticking out just below the 24-pin power; those are incredibly interesting, mildly frustrating, and something we didn't notice until later in the build. We'll cover that in a bit.

IN SHE GOES

With our brackets removed, we then added the standoffs for our Kraken cooler. It's really simple this time around—grab the AM5 threaded standoffs from the AM5 bag, screw them into position in the backplate, and that's it—super simple. We've also installed the memory here. In hindsight (particularly as we'd already built in this chassis once with this exact kit), we recommend you go with a lower-profile memory kit rather than a set of Dominator Titaniums. Sure, the speed is great and the performance stellar, but clearance is everything in a build like this.

To install them, open the latches at the end of each DIMM slot, line the stick up with the offset notch in the center, and push it into place until the latch locks down. Fortunately, as we're utilizing an ITX board, there's no need to worry about what slots you need to use. The answer is simple: both of them. With that, we can get the board into the Era 2's skeleton.

Lay the Era 2 on its side, line up the X870I-I in there, and carefully place it over the motherboard stand-offs before installing with the included motherboard screws found in the accessories box. You may need to lift the PCIe passthrough cable out of the way, along with the power passthrough, to fully secure it in position. Here, you'll notice that it's going to be a pain getting that 8-pin EPS power in for the CPU, and you'd be right. Once your motherboard is secure, pop that PCIe riser cable in **[STEP 9]**.

With that done, we're going to try to get ahead of the game and start cable managing as much as we can. That means installing as much of the front I/O as possible; USB Type C, Type A, and power can all go in. This is mostly a preliminary measure; you'll likely end up re-cabling these about 15 times before the end, but it's the thought that counts. With early cable management done, it's time to get the PSU in there. Slide the PSU bracket out, then secure it to your PSU with the included screws. Attach the power passthrough cable, and slide it back onto its bracket before securing with a screw on the far left ledge **[STEP 10]**.

More cable management now, this time those bottom fans into the CPU fan slot on the motherboard. We've also gone ahead and ran the CPU power cable around the back of the board, underneath the PCIe slot on the other side of the chassis, and into the EPS header on the motherboard, weaving it around the fans. It's a bit of a catch 22, this chassis. On the one hand, the Revolt SFX is great, because it has short cables; on the other, because it has short cables, certain runs become more challenging as a result. We've also gone ahead and installed the 24-pin power here too, weaving it and bunching it up in a way not to reduce total cable bulk trailing around the chassis **[STEP 11]**.

Next up, and perhaps earlier than you might expect, we're getting that RTX 5080 installed. Nice and simple, this, but we've

16

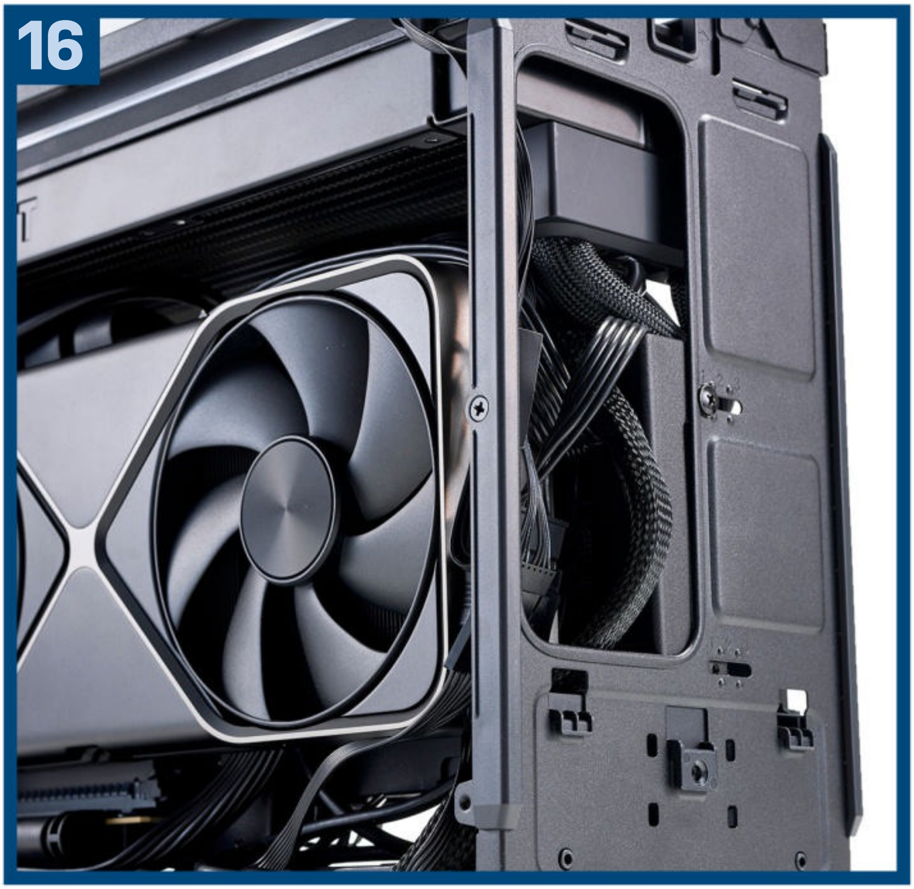

got a ton of cable management to do later for our 240mm AIO. To install the RTX 5080, line it up with the PCIe slot on the side, and push it into position until it clicks into place. Then, run the 12VHPWR into the angled connector on top, into the PSU, and you're all set **[STEP 12]**. Again, cable management is pivotal, so try to tuck cables and run them intelligently where you can, hide them out of the way, and use components to wedge them into position if you can. It's less about looking clean, and more about ensuring that we've got plenty of space for airflow and fewer conflicts with components (namely that cooler).

TINY AFTERTHOUGHTS

Small problems require big solutions. Maybe. Honestly, we're not entirely sure this is a necessary component, but here we are. **[STEP 13]** See that tiny chip board? That's included with the X870I as standard; it slots onto those two USB Type C ports we pointed out earlier, and points outward (toward the PSU), giving us access to two USB 2.0 headers, along with an HD Audio Passthrough port, plus two SATA headers.

The reason that's a critical component for us is because we need a 2.0 header to allow functionality with the NZXT pump and the motherboard (in fact, almost all AIO cooler solutions require a USB 2.0 header at this point, frustratingly).

Although it's useful, it feels like we could have gotten away with the two USB 2.0 headers and that HD audio passthrough on the board in the same space those USB Type C connectors occupy. Perhaps it's time to ditch the SATA headers, given how affordable and dense NAND flash has become on M.2 ports.

We foolishly, however, didn't realize that we needed this until after we'd begun to install the AIO. To prepare NZXT's latest Kraken coolers, line up the radiator, then the F240 Core frame, and then the bracket, and attach the three items together using the long included rad screws. Then feed the larger master cable into the controller in the radiator, and you can at last fit your CPU block on top of the CPU before routing all those cables. If you're utilizing an Era 2 like we are, plus this cooler, you're limited in fan orientation. Effectively, you're doing a pull solution, as if you do a push solution in anticipation of that master cable being a bit awkward to manage (with the RGB fans on the inside of the machine), it's likely that your fans will make contact with cables or other components, stopping them from spinning.

With CPU block attached (ideally with an extra helping of non-conductive thermal paste), you can then push the radiator bracket back on top of the chassis until it clicks and locks into position. Push the tubing up and next to the radiator. Then, like we did, scream internally as you realize that you don't have a 2.0 header, and go rummaging for the silly add-in board before finally fitting the cables into position **[STEP 14]**.

We're getting there. All that's left is to plug the pump into one of the fan headers, run the SATA power in, and tuck a few cables out of the way **[STEP 15]**. Plus, look at the size of that monstrous LCD block there. Shame it's going to be hidden behind that shell. Fractal, if you're listening: Era 3, slightly better cooling, and a window here, please.

Finally, with all the cables managed, we tucked everything out of the way around the side of the GPU, ensuring we tucked any spare front I/O cables and headers down along the side of the fans on the bottom. Cable management is your friend here; any cable ties or solutions you have will help a ton. SFX PSU is a must, any SFX-L units will make the build far more difficult, and it's just not worth the hassle **[STEP 16]**.

ACHIEVEMENT UNLOCKED

WE'VE REITERATED IT enough at this point, but it's worth mentioning again. This was a hard build. It might not seem it, but chassis like these tend to draw blood during the construction, even for our experienced hands. They're cramped, difficult to work in, and generally require considerably more thought, tact and expertize than what you'd find from a mid or full-tower chassis. That's not a bad thing, though. The Era 2 is incredible to look at, and performs fantastically throughout the whole procedure. You've got to bear in mind when you're looking at a chassis like this that it's built with this massive volume constraint placed on its head, and yet it still needs to fulfil the modern-day requirements that a good case needs to have. It needs to support modern hardware, modern specs, correct cooling, keep temperatures low, look the part, provide multiple build avenues, be available in different colorways, and most importantly, not break the bank for the company to manufacture and design. Getting all of that right is a massive balancing act.

With all that said, Fractal really nailed it on this one. Our only complaint is that perhaps it could do with a little more

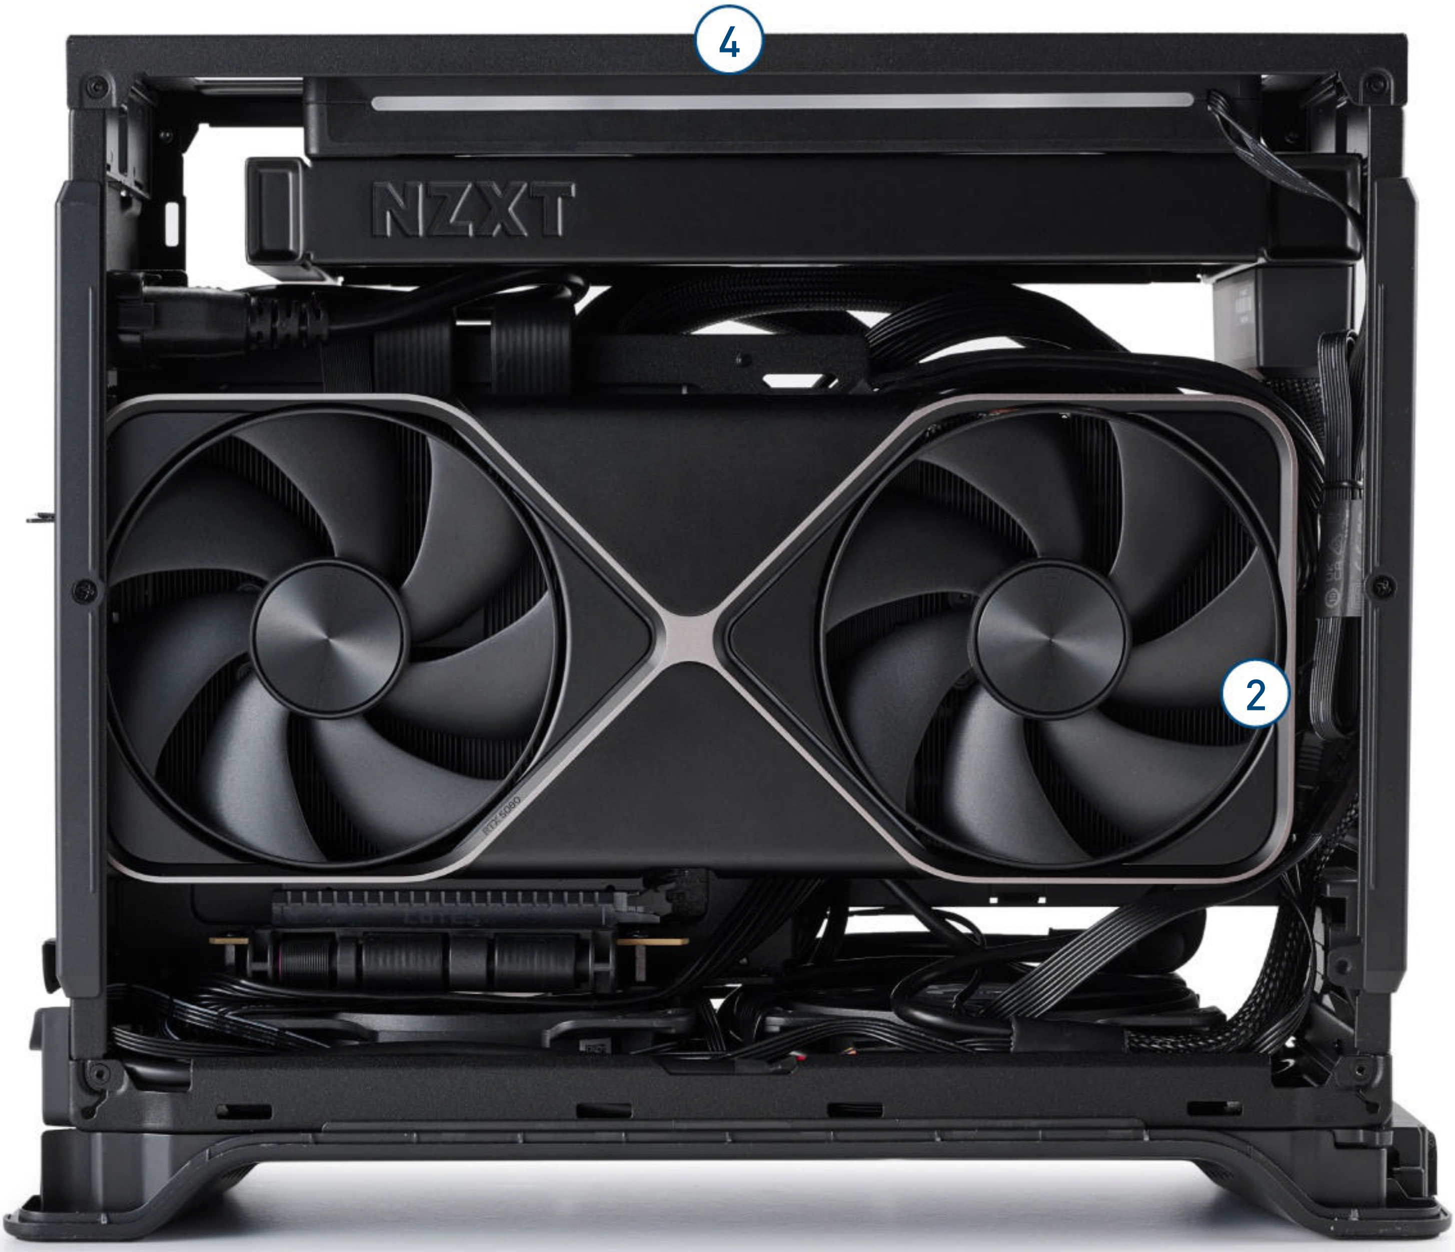

flexibility on that power supply mounting location, and it could probably do with a touch more clearance to allow for better cable management on the EPS/CPU power for the motherboard. And of course, the glass window comment. There's kind of an elegance in a lot of ways to that skeleton look as well. It's a bit niche, but that barebones look is pretty sick. We actually debated leaving the shell off entirely and doing something radical with our cover image, but that would take away from the satin paint job.

1 **Realistically opting for lower profile memory, and a slimmer AIO liquid cooler would make this build easier. Plus, we're effectively wasting cash, as you can't really see this when it's powered on, either.**

2 **If you want, you can grab a 1000W SFX PSU from Corsair and theoretically run an RTX 5090 in here as well—if you can find one in stock.**

3 **This motherboard is massively overkill for what we're doing here, but mostly on price. $520 for an ITX board is a bit tough even for us to swallow.**

4 **As you can probably tell, given the external shell is off as we take those photos, we do wish there was a glass panel version of the Era so we could show off more of that interior.**

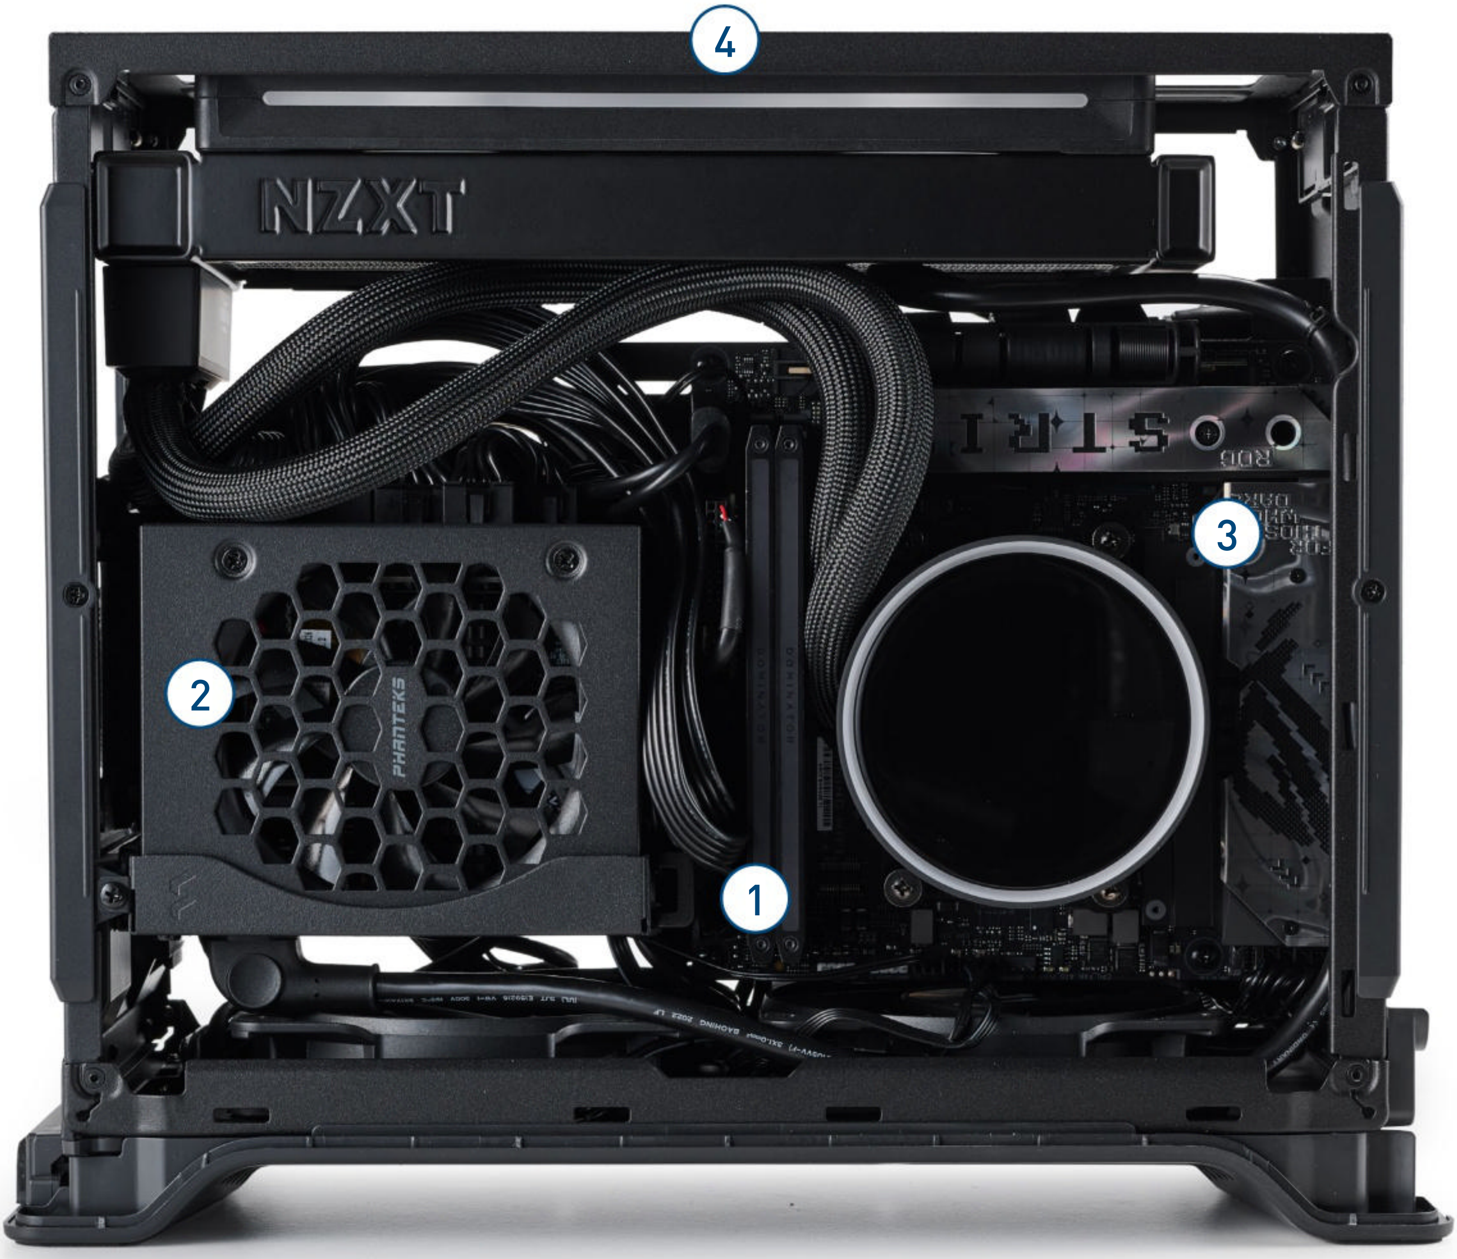

MICRO **MACHINE**

HOW WE TEST

Numbers, numbers, and more numbers. Analyzing data is everything to a good tech nerd, and we're no different here at *Maximum PC*. Understanding how each and every system we build performs is pivotal for us to give you a good critique of the system at hand and recommend whether you should pick up the parts to create your own.

Every system we finish and complete here in the *Maximum PC* lab is heavily tested across all manner of benchmarks and scenarios. We cover absolutely everything from CPU performance and SSD speeds to power draw, thermal performance, AI potency, and of course, gaming at both 1080p and 4K. There are a total of 98 test runs conducted, giving us a whole host of data points to draw from.

THIS ISSUE'S ZERO POINT

We're using exactly the same zero point as the previous issue. That's a custom RTX 4080 build, built off the back of Intel's 14th Gen 14900K under its revised microcode, and backed up by the RTX 4080 Gigabyte Aero OC (that equally performs identically to its Super cousin, too). Interestingly, the core pricing for our RTX 5080 build here is significantly lower than that little number, although it's worth pointing out that neither GPU is currently available for purchase, sadly. The RTX 4080 Aero OC holds a special position, as although slightly more expensive than its RTX 4080 Super successor, it performs just as well, and effectively bridges the gap between the three generations of GPU, giving a really clean look at just how well the RTX 5080 does in game.

Our Zero Point features the Intel Core i9-14900K, an Asus Z790 ROG Dark Hero, a Phanteks 360mm Glacier One 360D30 RGB AIO cooler, 32GB (2x16GB) of Lexar Thor OC DDR5 @ 6000 C32, a 2TB Crucial T700 M.2 PCIe 5.0 for the primary OS, a 2TB Kingston Fury Renegade PCIe 4.0 SSD for our secondary SSD, and a Gigabyte Aero OC RTX 4080 for our graphics card. All of this is housed nicely inside Geometric Future's M4 chassis, complete with 1200W Corsair RMx Shift PSU and a triple pack of Phanteks M25-120 fans.

The total price for the Zero Point build at time of writing is $3,499 (16.4 percent more expensive), with a core price of $2,682 (15 percent more expensive than our RTX 5080 build here).

GEN-ON-GEN PERFORMANCE BOOSTS

We've got to say, despite the ongoing madness around the ROP situation on these cards, and the lack of hot spot tracking, overall performance for the RTX 5080 is exactly where we'd expect from a generational jump-up: bang average. In our testing, 4K performance jumped up on average by around 27-percent, with 1080p being a little less dramatic at just a 13 percent increase on the average frame rates. It's respectable, but not massively groundbreaking, particularly as a lot of that performance boost is coming directly from titles that can take advantage of Nvidia's latest round of multi-layered frame generation.

That's particularly noticeable in *Cyberpunk*. At 4K, the difference with DLSS and Frame Gen across the two systems is astounding, going from 75.7fps on average to 103.9fps. It's no small jump, that's for sure. Still, there are some nice raw performance number bumps, too, including *Total War: Warhammer 3* shooting up from 67fps to 87fps too, and *Final Fantasy* bounced up quite significantly by a good 20fps here as well.

If you can get it, then the RTX 5080 is without a doubt the card we'd recommend for 4K gaming. It almost nails the price-to-performance metric—at least until we get hands on with the lower-specced cards.

1080P WOES ONCE AGAIN

Interestingly, we began to see 1080p bottlenecks again—or significant performance drops compared to its predecessor—in exactly the same manner as we saw with the RTX 5090. *Far Cry 6* frame rates slumped significantly, with the RTX 5080 landing a depressing 158.3fps on average frame rates versus the 4080's 171.3. It's hard to pin down what is causing this, as we've tested both 50 series cards on processors from Intel and AMD, suggesting it's not isolated to a platform, but something to do directly with the cards themselves, or the system. We're going to be doing some behind-the-scenes testing on this one to get to the bottom of it, and admittedly it's a bit of a niche scenario, but if you do play older games or unique titles, and the RTX 5080 seems to be chugging compared to older cards, it's certainly worth knowing about.

POWER DRAW AND THERMALS

The big elephant in the room has to be the power and heat, right? That's the de facto fear with any ITX build, and the Era 2 handles things nicely. Max temps topped out at 96.1 C on the CPU, and the RTX 5080 hit a top temp of 78 C—not bad for a twin-slot card in a tiny case. Those two base intake fans do some serious heavy lifting, and those CPU temperatures are within parameters there for an AMD chip at this caliber (it's not far off the temperatures it gets in our full-time testbed with a dedicated 360mm AIO and ten fans).

Power draw similarly is considerably lower than last time around, with the 5080 and Ryzen 9 drawing only 568.7W from the wall under load, giving us plenty of headroom with that 850W Revolt.

IS IT WORTH IT, THEN?

It's a tricky one. The numbers don't lie, and we've seen a significant boost in overall performance at both 4K and 1080p, far greater than what we saw with the initial RTX 40 Super series launch. It's not revolutionary by any length of the imagination, but it's enough that if you do decide to migrate from one generation to another, you won't be too disappointed.

The big advantage comes in the form of that extra frame-generation tech. If you're not susceptible to the latency or any artifacting, then it will undoubtedly take your 4K gaming to a whole other level.

The biggest problem though, is still stock. None of these cards are available right now. The only ones out there are either sold by unscrupulous retailers at $600 more than their launch price, or scalpers for near RTX 5090 retail prices. That's not something we can abide by really. Until Nvidia gets these things back in stock, at a reasonable volume, at a retail price, it's a real struggle to recommend.

SYSTEM BENCHMARKS

For our system testing, we use a multitude of separate benchmarks. For our CPU tests, these include Cinebench 2024 and Geekbench 6.3.0. All results are indexes in this case. Multicore per thread is achieved by dividing the total multicore scores by the number of threads on the CPU. For storage testing, we're taking advantage of CrystalDiskMark 8.0.5—the best scores are in bold, and all results are reported in MB/s. Throughout our entire benchmarking process, we test each system with a power meter plug running, and the latest version of HWMonitor installed and up as well. This allows us to report on the maximum power draw and temperatures seen over the entire benchmarking process.

CPU PERFORMANCE CINEBENCH 2024 & GEEKBENCH

	ZERO-POINT	
Cinebench 2024 Multi-Core	1,923	1,745 (-9.3%)
Cinebench 2024 Multi-Core Per Thread	60	73 (21.7%)
Cinebench 2024 Single-Core	125	133 (6.4%)
Geekbench 6 Multi-Core	20,016	19,953 (-0.3%)
Geekbench 6 Multi-Core Per Thread	625.50	831.38 (32.9%)
Geekbench 6 Single-Core	2,997	3,323 (10.9%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

STORAGE PERFORMANCE CRYSTALDISKMARK 8.0.5

	ZERO-POINT	
C: Sequential Read/Write Q8	**11,823 / 11,275**	10,252 / 8,678
C: Random 4K Read/Write Q1	63.04 / 301.83	**76.05 / 312.25**
D: Sequential Read/Write Q8	6,839 / 6,659	N/A
D: Random 4K Read/Write Q1	89.38 / 329.38	N/A

POWER DRAW & TEMPERATURE HWMONITOR & DEDICATED POWER METER

	ZERO-POINT	
Max CPU Temperature	100.0 C	96.1 C (3.9%)
Max GPU Temperature	72.0 C	78.0 C (-.8.3%)
Max System Power Draw	618.2 W	568.7 W (8.0%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

GAMING BENCHMARKS

Game tests are performed on Ray Tracing Ultra *(Cyberpunk)*, Ultra *(Far Cry 6)*, Maximum *(Final Fantasy: XIV Dawntrail)*, and Ultra *(Total War: Warhammer 3)* graphics presets. Three runs are performed per title, per resolution. *Cyberpunk* is always tested with ray tracing enabled regardless. Under its second set of conditions, DLSS/FSR/XeSS is enabled in *Cyberpunk* dependent on GPU manufacturer, with upscaling preset set Quality. Nvidia cards also have AI Frame Generation enabled. *FF XIV: Dawntrail* is utilizing FSR across the board regardless of GPU, with LOD on Distant Objects enabled, and dynamic resolution disabled. *Far Cry 6* is tested with DXR disabled, and the HD Texture pack downloaded and enabled.

1080P PERFORMANCE

	ZERO-POINT	
Cyberpunk 2077 Stock (Avg/Min FPS)	85.2 / 68.9	**99.3** / **86.4**
Cyberpunk 2077 AI Upscaling (Avg/Min FPS)	202.4 / 157.9	**228.7** / **176.5**
Far Cry 6 (Avg/Min FPS)	**171.3** / **134.0**	158.3 / 117.3
Final Fantasy XIV: Dawntrail (Avg/Min FPS)	208.5 / **89.0**	**241.8** / 84.3
Total War: Warhammer 3 (Avg/Min FPS)	161.4 / 111.7	**204.2** / **147.7**

1080P PERFORMANCE

	ZERO-POINT	
1080p Titles Aggregate (Average FPS)	165.77	186.47 (12.5%)
1080p Titles Aggregate (Minimum FPS)	112.30	122.45 (9.0%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

4K PERFORMANCE

	ZERO-POINT	
Cyberpunk 2077 Stock (Avg/Min FPS)	26.7 / 17.6	**34.0** / **30.3**
Cyberpunk 2077 AI Upscaling (Avg/Min FPS)	75.7 / 39.2	**103.9** / **95.2**
Far Cry 6 (Avg/Min FPS)	108.3 / 76.7	**131.7**/ **116.7**
Final Fantasy XIV: Dawntrail (Avg/Min FPS)	94.6 / 57.3	**115.1** / **77.0**
Total War: Warhammer 3 (Avg/Min FPS)	66.8 / 50.7	**86.9**/ **66.0**

4K PERFORMANCE

	ZERO-POINT	
4K Titles Aggregate (Average FPS)	74.42	94.30 (26.7%)
4K Titles Aggregate (Minimum FPS)	48.30	77.03 (59.5%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

SYNTHETIC GAMING BENCHMARKS

We're using 3DMark's latest Speed Way, Steel Nomad, Night Raid, and CPU Profile tests. Speed Way represents a DX12 title at 1440p, with ray-tracing elements, Steel Nomad is DX12 4K without ray tracing, and Night Raid (DX12) is designed with integrated GPUs in mind at 1080p. CPU Profile acts as an additional CPU benchmark (reporting on max thread output only), with specific gaming workloads prioritized. Each synthetic test is run three times, and indexes are reported on.

3DMARK PERFORMANCE

	ZERO-POINT	
Speed Way (DX12 + RTX @ 1440p)	7,246	8,951 (23.5%)
Steel Nomad (DX12 @ 4K)	6,541	8,154 (24.7%)
Night Raid (DX12 @ 1080p)	95,238	93,992 (-1.3%)
CPU Profile	15,312	13,597 (-11.2%)
		0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

AI INFERENCE BENCHMARKS

For our AI testing, we use UL Procyon's Professional Benchmark Suite. We focus on its AI Computer Vision (which gives us a gauge on how our systems operate when dealing with inference engine tasks), and AI Image Generation benchmarks. For Computer Vision, we report on the total inferences performed, the average inference time, plus the Index score across all tests. The Index score is more valuable, as it's the time taken between inferences that indicates performance. We run all tests on AI Vision using Float16 precision. We test both the CPU and the GPU under Microsoft's ML API.

For Image Generation, we test under Stable Diffusion 1.5 (Float16 precision), using all engines available. However for comparison, we only include scores from the same engine. We report on the Index score, the time taken, and the image generation speed. All results are averaged across all tests. Best scores, where appropriate, are in bold.

AI COMPUTER VISION PERFORMANCE

	ZERO-POINT	
Microsoft ML API CPU (Index)	50	73 (46.0%)
Microsoft ML API GPU (Index)	1,902	2,324 (22.1%)
		0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

	ZERO-POINT	
Microsoft ML API CPU (Total Inferences)	10,284	14,047
Microsoft ML API CPU (Average Inference Time)	659.88 ms	**401.14** ms
Microsoft ML API GPU (Total Inferences)	226,692	251,215
Microsoft ML API GPU (Average Inference Time)	9.883 ms	**7.50** ms

AI IMAGE GENERATION PERFORMANCE

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Index)	2,616	3,464 (32.4%)
		0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

	ZERO-POINT	
Stable Diffusion 1.5 ONNX (Time Taken)	38.223 seconds	**28.865** seconds
Stable Diffusion 1.5 ONNX (Image Generation Speed)	2.389 seconds/image	**1.804** seconds/image

NVIDIA BLACKWELL RTX DEEP DIVE

The biggest changes all revolve around AI

IT TOOK A BIT LONGER than we were expecting, but Nvidia's Blackwell RTX architecture and GPU family have finally arrived. We reviewed the RTX 5090 last issue, and test the RTX 5080 in this issue, but there's a lot more to say about everything that's going on under the hood.

Blackwell RTX differs in several key areas from the data center Blackwell architecture, just as the Hopper architecture differed from Ada Lovelace design. As expected, AI and machine learning lands front and center for all the new GPUs coming out this year, which can be both good and bad.

Beside the core architectural changes relative to the prior generation, Nvidia has overhauled its Founders Edition design, created new software tools and features, doubled down on frame generation as a way to 'boost' performance, and added new functionality that can help with professional workloads.

The RTX 50-series family of GPUs will carry the torch for the next two years of Nvidia graphics cards, give or take. Some of the technology is forward-looking, and could take at least that long before it starts to see significant uptake by games and applications, and other features will have an almost immediate impact. So put on your thinking cap, and let's discuss the ups and downs of the new architecture.

The RTX 5090 Founders Edition packs a whopping amount of compute and power into an impressively thin 2-slot card.

BLACKWELL RTX SPECIFICATIONS

The core specifications for the Blackwell RTX GPUs look like the prior-generation Ada Lovelace architecture, outside of the top RTX 5090. GB202 represents a sizeable upgrade from AD102—literally. It's 750 mm^2 compared to 608 mm^2—a 23 percent increase in die area. It also has up to 192 SMs, compared to a maximum of 144 SMs—a 33 percent increase that mostly overlaps with the die size change. Below GB202, however, everything seems to be far less impressive in terms of gen-on-gen upgrades.

GB203 tops out at 84 SMs (the same as AD103), with the RTX 5080 using the fully enabled chip, while the RTX 5070 Ti trims things down to 70 active SMs. GB205 drops down to just 50 SMs at most, 46 of which are enabled on the RTX 5070—the same number as used in the RTX 4070, but AD104 offered up to 60 SMs.

There are supposed to be two more chips: GB206 for the RTX 5060, and 5060 Ti (including laptop variants) with up to 36 SMs, and GB207 for RTX 5050 (desktop and laptop), with up to 20 SMs. It will be interesting to see if Nvidia actually launches a desktop 5050, considering it skipped the desktop RTX 4050 on the previous generation.

But across the GPU stack, excepting the biggest chip, the new Blackwell chips are equal to or smaller than their predecessors. Does that mean we shouldn't expect much in the way of performance improvements? The answer ends up being quite nebulous, as it depends heavily on how you want to quantify performance. Let's delve deeper into what makes the Blackwell RTX architecture tick.

BLACKWELL'S AI ASPIRATIONS

At the heart of everything Blackwell, Nvidia has made some key changes. Superficially, it's easy to look at the performance and specifications, and conclude that Blackwell is merely Ada version 2.0—or not even that (see TSMC 4N sidebar). But despite a lot of similarities, the two architectures are definitely not the same. Blackwell builds on Ada, just as Ada built on Ampere, and so on back to Nvidia's first GPUs. Here's an overview of the changes.

AI enhancements are a big deal, with the chief new addition being native support for FP4 and FP6 number formats on the tensor cores. FP4 requires half the memory space, and offers double the computational throughput compared to FP8. 4-bit floating point versus 8-bit floating point will do that, just as FP8 used half the storage and offered twice the throughput of FP16.

The lack of numerical precision needed for AI workloads has been a source of continual surprise. Scientific workloads still generally rely on FP64 for accuracy, but machine learning has found ways to trim that down to FP16 almost from the start of the current era—if we classify everything in the post-AlexNet time frame (post 2012) as 'current'. Then, things moved to 8-bit using both integers and floating-point formats, and now there's active use of just 4-bit formats.

FP4 seems a bit crazy. With only 16 possible states, on the surface, it doesn't look particularly useful. But the format uses a sign bit with three exponent bits, with a base 4 format, giving a range of ±64 (±43). It can also be used to represent numbers as small as ±4-3. The net result is that a combination of clever ultra-low precision techniques can result in similar accuracy to FP8, and even FP16 formats (within ~2 percent), though it's not clear if these techniques are universally applicable, or only function with a certain subset of workloads.

Nvidia classifies the FP4 performance as 'TOPS', or TeraOPS (trillions of operations per second), even though normally TOPS refers to integer calculations, while TFLOPS is used for floating point numbers. But the industry at large has been moving toward

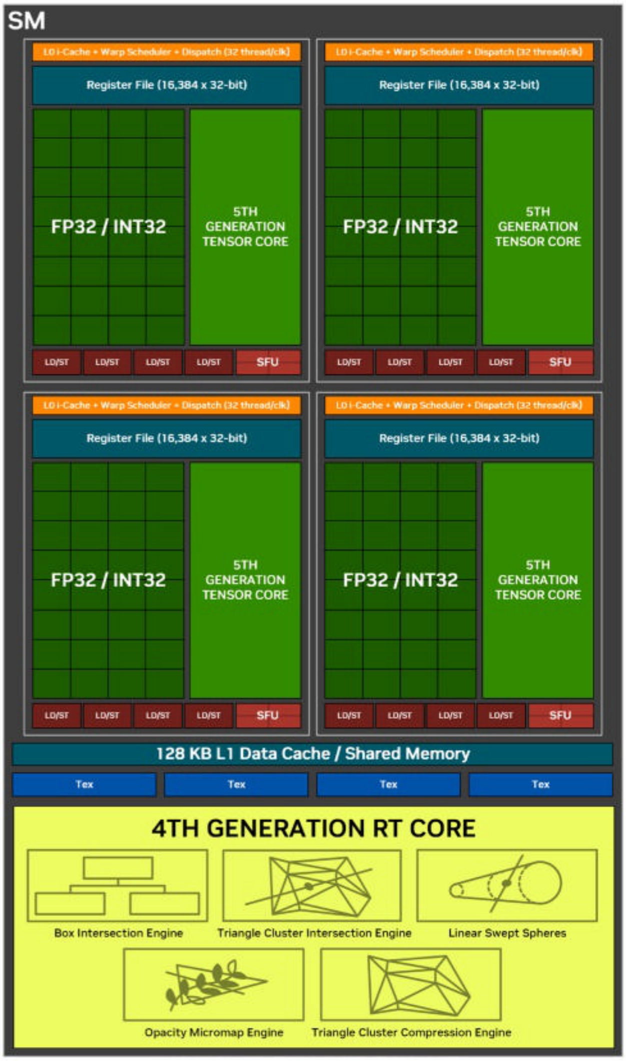
SM
L0 i-Cache + Warp Scheduler + Dispatch (32 thread/clk)
Register File (16,384 x 32-bit)
FP32 / INT32
5TH GENERATION TENSOR CORE
LD/ST
SFU
128 KB L1 Data Cache / Shared Memory
Tex
4TH GENERATION RT CORE
Box Intersection Engine
Triangle Cluster Intersection Engine
Linear Swept Spheres
Opacity Micromap Engine
Triangle Cluster Compression Engine

The Blackwell SM has been upgraded, with double the INT32 shaders and 5th generation Tensor cores that support FP4.

The GB202 is Nvidia's largest 'consumer' chip since the RTX 2080 Ti, and the first 512-bit memory interface since the GTX 580.

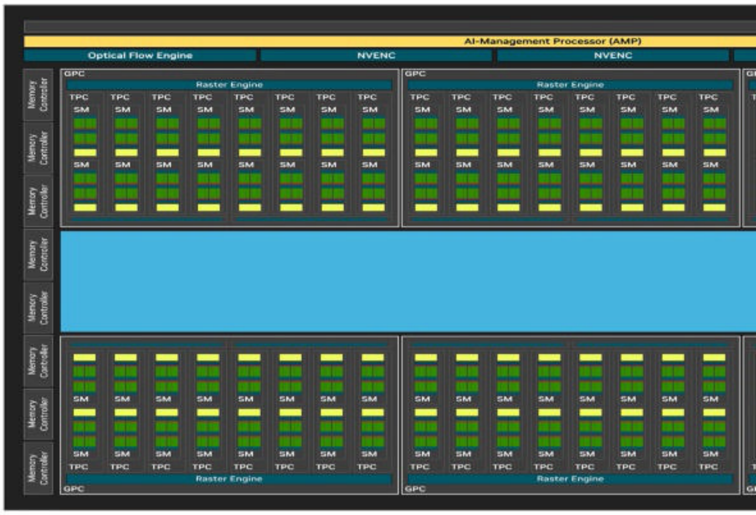
AI-Management Processor (AMP)
Optical Flow Engine
NVENC
NVENC
GPC
Raster Engine
TPC
SM
Memory Controller

GRAPHICS CARDS

	RTX 5090	RTX 5080	RTX 5070Ti	RTX 5070
Architecture	GB202	GB203	GB203	GB205
Process Technology	TSMC 4N	TSMC 4N 00	TSMC 4N	TSMC 4N
Transistors (Billion)	92.2	45.6	45.6	31
Die size (mm^2)	750	378	378	263
SMs	170	84	70	48
GPU Shaders (ALUs)	21,760	10,752	8,960	6,144
Tensor / AI Cores	680	336	280	192
Ray Tracing Cores	170	84	70	48
Boost Clock (MHz)	2,407	2,617	2,452	2,512
VRAM Speed (Gbps)	28	30	28	28
VRAM (GB)	32	16	16	12
VRAM Bus Width	512	256	256	192
L2 Cache	96	64	48	48
Render Output Units	176	112	96	80
Texture Mapping Units	680	336	280	192
TFLOPS FP32 (Boost)	104.8	56.3	43.9	30.9
TFLOPS FP16 (FP4 TOPS)	**838** (3,352)	**450** (1,801)	**352** (1,406)	**247** (988)
Bandwidth (GB/s)	1,792	960	896	672
TBP (watts)	575	360	300	250
Launch Price	$1,999	$999	$749	$549

© NVIDIA

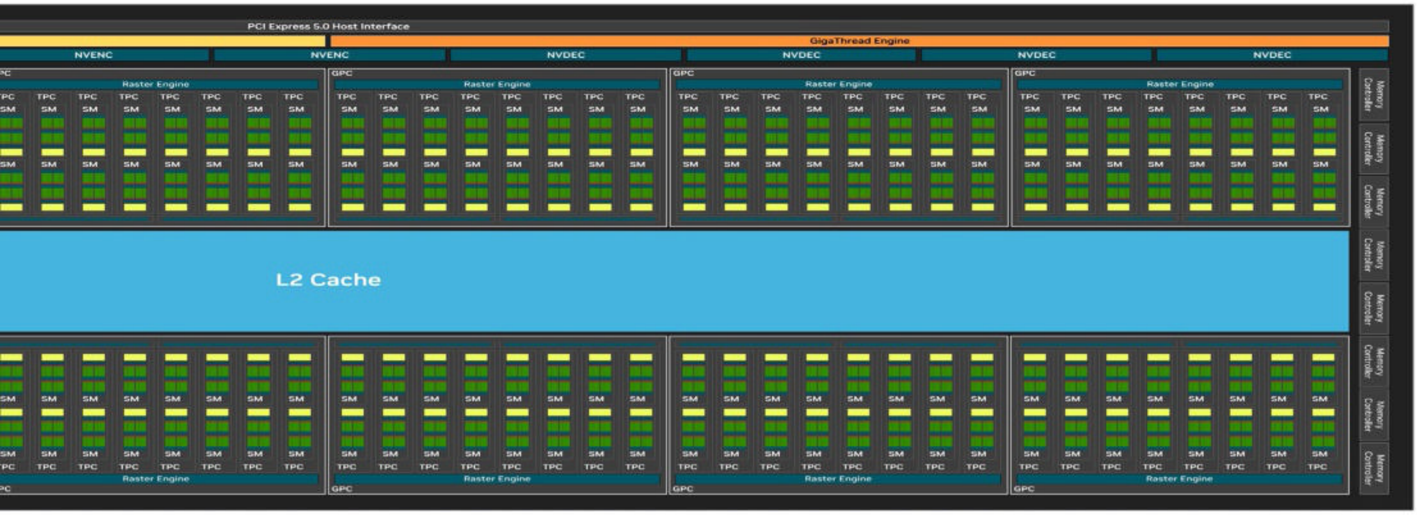
PCI Express 5.0 Host Interface
GigaThread Engine
NVENC
NVENC
NVDEC
NVDEC
NVDEC
NVDEC
GPC
Raster Engine
TPC
SM
L2 Cache
Memory Controller

NVIDIA ARCHITECTURES PAST TO PRESENT

Throughout this deep dive, we'll refer to prior-generation hardware from Nvidia, and it's important to understand what we're talking about. However, we're only going back as far as the first RTX cards, which launched in 2018—we'll leave out the earlier GTX and other generations. Here's the short rundown, from oldest to newest.

TURING (2018): The RTX 20-series arrived in 2018 and marked a big change in Nvidia's graphics hardware. It came with new branding (RTX instead of GTX), and included two major new hardware innovations: hardware ray tracing functionality, and AI-focused tensor cores. Each Streaming Multiprocessor (SM) contained 64 CUDA cores capable of two 32-bit FMA (fused multiply accumulate) operations per cycle, with eight tensor cores capable of 128 16-bit floating point (FP16) FMA operations per cycle. It was manufactured on TSMC's 12nm FinFET process node.

AMPERE (2020): RTX 30-series GPUs moved to Samsung 8N (a refinement of Samsung's 10nm node), and doubled down on several key aspects of the prior generation. Each SM now contained 64 CUDA cores that could do FP32 or INT32 operations—the same as before—with an additional 64 CUDA cores that could only do FP32 instructions. That doubled the graphics throughput. The number of tensor cores per SM was reduced to four, but each one is twice as fast for baseline performance—256 FP16 operations per cycle. The tensor cores also add sparsity support, which again doubles the potential compute for 'sparse' matrices—basically, it can skip half of the multiply by zero calculations (as the result is always zero). The ray/triangle intersection hardware also became twice as fast.

ADA LOVELACE (2022): RTX 40-series shifted manufacturing back to TSMC, using a tuned 4N ('4nm Nvidia') node that was an optimization of its existing 5nm-class process. That brought big improvements in efficiency. Clock speeds also saw a jump, and the ray/triangle intersection rates doubled. Each SM still has 128 FP32 CUDA cores, 64 of which support INT32, and four tensor cores. The tensor cores still perform up to 256 FP16 operations per cycle—double with sparsity—but they can also execute 512 FP8 operations per cycle.

TOPS, particularly with regards to things like NPUs (Neural Processing Units), and Nvidia seemingly wants to make it clear just how much more potential performance its Blackwell GPUs offer. The RTX 5090, for example, delivers up to 3,352 theoretical TOPS of FP4 compute.

ENHANCING THE STREAMING MULTIPROCESSORS

Blackwell changes the SM and CUDA core arrangement (see 'Past to Present' sidebar). There are now 128 'full citizen' CUDA cores per SM, each capable of both FP32 and INT32 operations. That's the same FP32 as the past two architectures, but double the INT32 rate. Nvidia's reason for the doubling of the INT32 capabilities traces back to AI yet again. INT32 gets used for a lot of memory address calculations, and AI needs that more than typical graphics workloads.

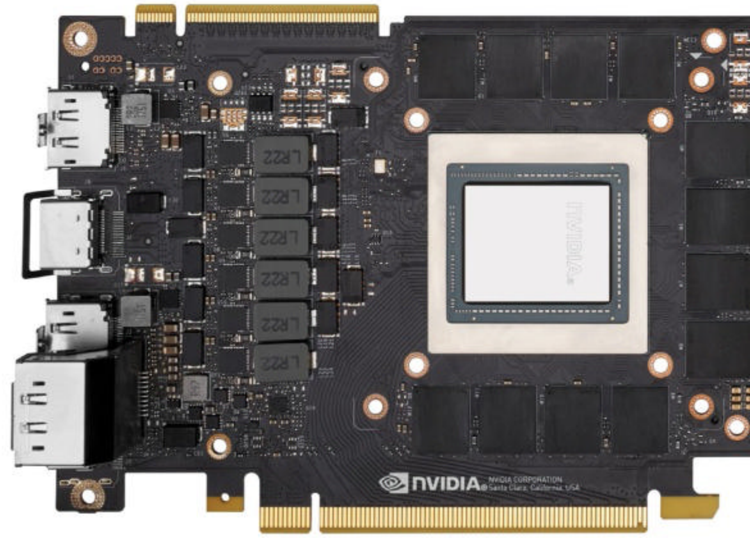

© NVIDIA

Ray tracing hardware also gets another bump, with double the ray/triangle intersections as before—eight times the rate of the Turing architecture. There are also enhancements to SER (Shader Execution Reordering), first introduced with Ada, that help it improve throughput in heavy ray tracing workloads. SER is now twice as fast.

The other biggest architectural change relates to 'neural rendering', which Nvidia classifies as the use of AI in the rendering pipeline. Prior architectures were more limited in the use of tensor operations alongside shader code, and Blackwell has altered some things to allow direct execution of tensor operations within shaders.

This should prove beneficial for things like Neural Texture Compression and Neural Materials, among other uses. What's not clear is how this relates to earlier Nvidia GPUs—can they accomplish the same tasks, just slower? Or is this new functionality that will only work on Blackwell and later GPUs? The former seems to be the case, but it's also something that appears to require direct developer adoption, so it could be a while before this sort of use gains traction.

MANAGING RESOURCES

With the increased emphasis on AI workloads, intermixed with other tasks like graphics work, comes the need to better manage the scheduling of resources. GPUs can now run a mixture of traditional graphics, ray tracing, and AI workloads simultaneously, with varying requirements for responsiveness.

Take a hypothetical future game that has AI-powered NPCs, ray-traced graphics, and DLSS upscaling and framegen. If the player interacts with an NPC by speaking to it, they won't want a massive delay before the NPC starts responding. At the same time, having graphical stutters while the AI begins processing what was said and how to respond could be just as jarring.

Blackwell adds a new AI Management Processor (AMP) to help keep things running smoothly. For example, in the case of a text-generation LLM, the time it takes for the LLM to begin responding can be more critical than the subsequent speed of text generation. The AMP would prioritize the time to first token in this case, but then scale back the amount of resources to a reasonable speed of continued response.

How this plays out will depend on the capabilities of the GPU being used as well as the task at hand. An RTX 5090, as

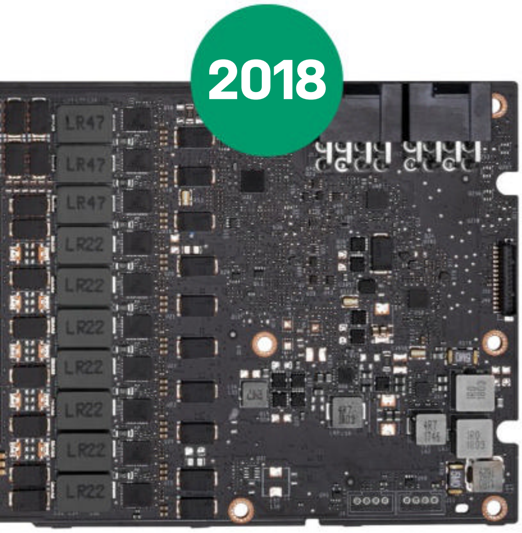

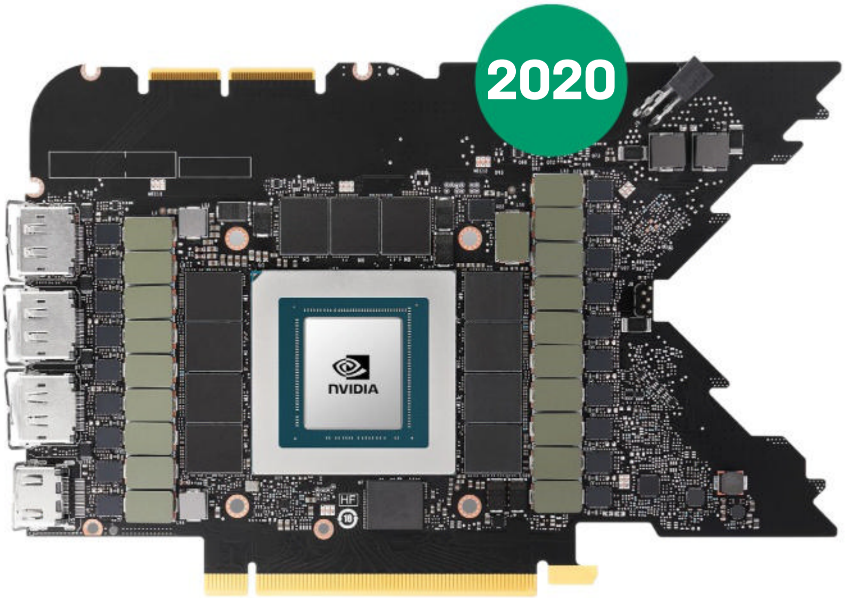

The RTX 2080 Ti FE (left) used a full-length PCB, and to improve cooling Nvidia cut out a notch at the back of the 3090 PCB (right). The 4090 GPU (bottom) kept the same core design but nearly tripled the transistor count of its predecessor.

an example, should be able to juggle more tasks while keeping gaming frame rates high, while an RTX 5060 may struggle to keep up if too many things are happening at once. We've already seen situations where the RTX 4060 can fall on its face—4K ultra with ray tracing, upscaling, and frame generation often proves to be too much for the little GPU—but hopefully the AMP can mitigate such situations.

MEMORY UPGRADES AND SIDEGRADES

Graphics memory has become a critical aspect of GPUs in recent years. Not only are games using increasing amounts of memory, but LLMs can be particularly voracious—that's why they're called large language models. It's not just a question of capacity, but also bandwidth.

Blackwell tackles the latter quite nicely, thanks to the use of GDDR7 memory. So far, the RTX 5080 gets 30 Gbps memory, while the other three announced GPUs will use chips clocked at 28 Gbps. That's a 40 percent increase in bandwidth for the 5070-class compared to the 4070-class' 20 Gbps GDDR6X, and a 30 percent improvement for the 5080 compared to the 4080 Super—or 34 percent higher relative to the original 4080.

And then there's the RTX 5090. It has the same 28 Gbps memory clocks as the 5070-class GPUs, but it also has a significantly wider 512-bit interface. That's double the width of the 5080 and 5070 Ti. The last time Nvidia did a true 512-bit memory interface for a GeForce card was with the GTX 280 and 285 all the way back

TSMC 4N RIDES AGAIN

One of the expectations from Blackwell was that it would move to a newer process node in order to improve transistor density, performance, and efficiency. Nvidia already revealed the GB200 GPU design for its data center parts, which moved to a refined 4NP (4nm Nvidia Performance) node, and it was rumored that the consumer GPUs would make the same transition. They didn't.

Blackwell RTX GPUs are manufactured on the same TSMC 4N node as the Ada generation. This has led to many referring to the new GPUs as a minor upgrade, like 'Ada v1.1' or calling the RTX 5090 the 'RTX 4090 Ti'. That's not entirely fair, as the rest of our deep dive shows how many changes have occurred, but it's also not entirely unfair.

We also don't know precisely what changes Nvidia and TSMC made with '4N' compared to the standard N4 node, nor how 4NP compares to either 4N or N4P. One of the supposed changes is that the number of metal layers was increased relative to N4P, potentially to improve clock speeds.

But all of these nodes are really variants and enhancements of TSMC's 5nm N5 process. N4 supposedly offers a five percent improvement in speed, or a 10 percent reduction in power. N4P was claimed to deliver an additional six percent higher performance, 11 percent higher efficiency, and a six percent increase in transistor density—along with a reduction in mask counts. What that means for the final processors using 4N isn't entirely clear.

What we do know is that the variations are all considered part of the 5nm-class family of nodes. There are benefits and drawbacks to each, but none of them offer the full-step improvement that was seen, for example, when looking at N5 versus N7. We'll need to move to a 3nm-class node for bigger density, performance, and efficiency gains, which likely won't happen until Nvidia's future RTX 60-series architecture.

MOBILE GPU SHENANIGANS

We've primarily focused on the desktop Blackwell GPUs, but Nvidia will naturally also have laptop variants. It has even announced the names and a few of the specifications, even though the actual laptops are still a month or two out. The various models are basically taking the same 'desktop GPU minus 10' approach as the RTX 40-series.

The top mobile GPU will be the RTX 5090, with 82 SMs and 10496 CUDA cores. That's basically the same as the desktop 5080, minus two SMs, and the 1824 'AI TOPS' indicates boost clock speeds should also be similar. Real-world clocks, with a TGP of 95–150 watts, will likely be lower. But what's most interesting is that the 5090 laptop GPU will have 24GB of memory on a 256-bit interface, meaning it has 24Gb (3GB) chips, unlike any of the other announced Blackwell GPUs.

The RTX 5080 laptop GPU lands below the desktop 5070 Ti, with 60 SMs and 7680 CUDA cores. It still has at least 16GB of memory, with up to 1334 AI TOPS, but again, power limits of 80–150 watts suggest it will perform worse in practice.

The laptop 5070 Ti continues the pattern, looking the same as the desktop 5070. It will have 46 SMs and 5888 CUDA cores, with 12GB of VRAM—slightly less than the desktop variant. Power limits will range from 65–115 watts, depending on the laptop.

The mobile RTX 5070 is a completely different chip. It will have 36 SMs and 4608 CUDA cores, with a 50–100 watts power limit. The GPU gets paired with 8GB of GDDR7 and a 128-bit interface. That's unfortunate from a memory capacity viewpoint. It also suggests that Nvidia will stick with 8GB of VRAM on the desktop 5060, with rumors of a desktop 5060 Ti with a 128-bit interface and either 16GB or 8GB of memory—just like the previous RTX 4060 Ti.

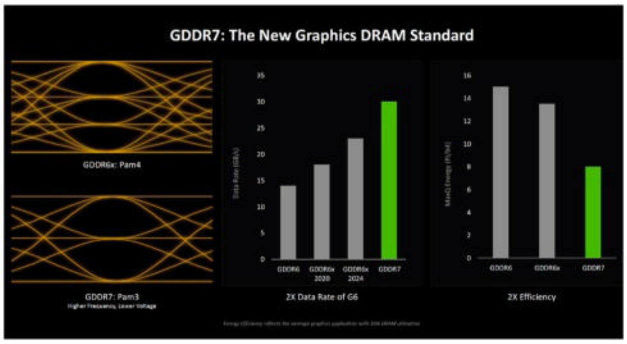

GDDR7 will boost memory speeds and bandwidth by up to 60 percent compared to GDDR6 and 30 percent over GDDR6X.

in 2008/09. A 33 percent wider bus coupled with 33 faster chips yields a 78 percent increase in memory bandwidth, compared to the previous-generation RTX 4090.

All of the currently announced RTX 50-series GPUs make use of GDDR7 memory, and that includes both desktop and laptop GPUs. Considering the not-yet-announced RTX 5060-class desktop parts will almost certainly use the same core hardware as the laptop RTX 5070 GPU (see Mobile GPU Shenanigans sidebar), it's a safe bet that Blackwell only supports GDDR7 memory.

What we didn't get—yet—is support for higher memory capacities on the existing bus widths. We know that 24Gb (3GB) GDDR7 chips are in development—in fact, they're being used already in the laptop 5090—but so far, none of the desktop GPUs are using anything other than 16Gb (2GB) chips. That means we'll have the same 16GB, 12GB, and (eventually) 8GB cards as the 40-series. Sigh.

Check back next year, and we'd say it's a very safe bet that we'll get a mid-cycle RTX 50-series refresh where many or all of the new variants will be upgraded to 3GB chips, with 50 percent more VRAM capacity. If you're not in a rush to upgrade, holding out for those future 'Super' models might be a reasonable plan.

MARKETING PERFORMANCE WITH MFG

Perhaps more than any other aspect of Blackwell, real-world performance will depend heavily on what exactly gets measured. This is where Nvidia's marketing department seems to be working overtime to make the new chips sound significantly better than their predecessors. Most of the architectural changes discussed so far seem relatively tame—doubling AI performance by using FP4 instead of FP8, for example. But just like the RTX 40-series added a new feature with frame generation, Blackwell and the RTX 50-series have added Multi Frame Generation (MFG).

If you're not a fan of framegen, we can't really blame you. Generated frames inherently feel different to rendered frames. If you can push the frame rate high enough, perhaps it doesn't matter so much, but more modest frame rates have proven particularly problematic for framegen. With MFG, things get even worse in talking about 'performance'.

Consider the case of 100fps, with and without frame generation and MFG. For traditional rendering, 100fps means the whole world state, including user input, gets updated every 10 ms. The result is an extremely responsive game. In the past, Nvidia pushed this narrative with its 'Frames win games' marketing that aimed to show how 240Hz, and even 360Hz and 480Hz displays,

Check back next year, and we'd say it's a very safe bet that we'll get a mid-cycle RTX 50-series refresh where many or all of the new variants will be upgraded to 3GB chips, with 50 percent more VRAM capacity

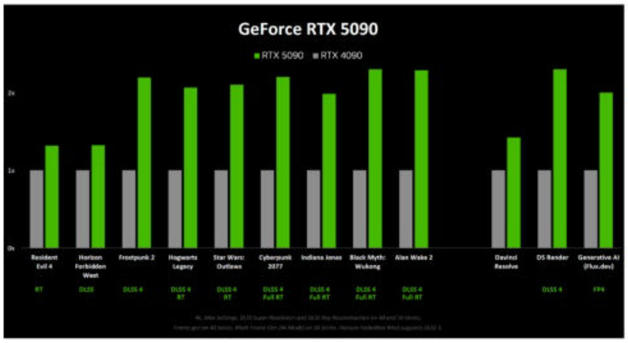

Nvidia's performance claims rely on MFG and FP4. The two left games don't use MFG, and show more modest generational gains.

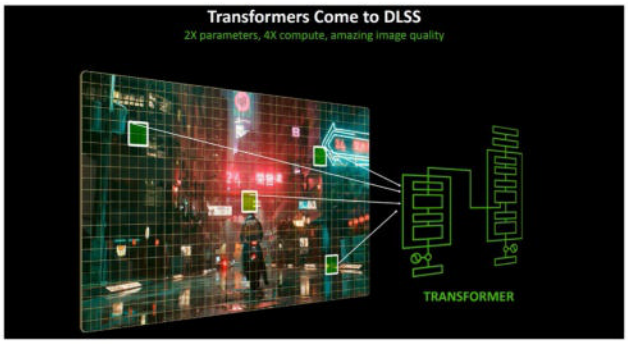
Transformers Come to DLSS
2X parameters, 4X compute, amazing image quality
TRANSFORMER

Transformers let AI models better focus on data. DLSS 4 has twice as many parameters as DLSS 2/3 and uses four times the compute.

were useful for competitive gaming. The core idea was to reduce latency and increase frame rates as much as possible, with a particular focus on lighter Esports titles.

With frame generation, 100fps means game world and user input updates actually happen at 50 samples per second, or once every 20 ms. Even worse, perhaps, is that the game needs to render two frames and then interpolate an intermediate frame, which delays monitor updates by two frames (relative to the framegen rate). Nvidia requires games with DLSS 3 framegen support to also support Reflex, an input latency-reducing algorithm, to try to mitigate the impact of the added latency. If you compare Reflex latency without framegen to the framegen numbers, you always end up with a more responsive game without framegen.

Now, Blackwell has added MFG, which does what the name implies. Instead of interpolating a single frame between two rendered frames, it can generate one, two, or three—Nvidia calls these modes MFG2X, MFG3X, and MFG4X, because the generated frame rate is two to four times the base frame rate. MFG tends to run fast enough that the gains are almost linear, so MFG2X might run 1.8X faster than native, MFG3X could be 2.4X faster, and MFG4X would be 3.2X faster. But if you consider the final fps, in the case where you only get 100fps, a game would be sampling user input just 25 times per second, and updating the world every 40 ms. Put bluntly: it is not the same.

To an external observer, the three examples of 100fps gameplay may look nearly the same. To a person physically playing a game, however, there's no question that native 100fps feels different to 100fps with framegen, and that also feels different than 100fps with MFG4X. Nvidia marketing will point to frame rates that are nearly twice as fast as the RTX 40-series GPUs, thanks to MFG, but in practice, it often ends up being a smoother presentation of frames with the same input rate.

SUBJECTIVE AND OBJECTIVE GAMING PERFORMANCE

Having tested MFG, framegen, and native resolution gaming on several GPUs, the resulting performance ends up very misleading if we just quote the MFG frame rates. It's not that MFG can't look and feel smoother than non-MFG, but rather that the gains aren't as substantial as the numbers would suggest.

Doing some fuzzy approximations, take a game running at 40fps natively. With a perfect doubling via framegen pushing that to 80fps, in practice, we'd say it feels more like a 20 percent improvement. With MFG4X quadrupling that to 160fps, it feels like perhaps another 20 percent improvement over framegen. Subjectively, then, it would be like 40, 48, and 58fps, while the frame rate to the monitor would be 40, 80, and 160.

But you won't normally get a perfect doubling or quadrupling of frame rates, and that word 'feels' does the heavy lifting. Many can and do argue for different evaluations of how much framegen and MFG help, and they're entirely within their rights. Just as evaluating how fast a game needs to run to feel 'smooth' can be very subjective, evaluating framegen and MFG is, if anything, even more subjective.

What happens if we just wipe MFG and framegen out of the picture, and look at the raw performance improvements from the Blackwell RTX GPUs? Not surprisingly, the gains are far more muted. The RTX 5090 was up to 40 percent faster than an RTX

NEURAL RENDERING

Beyond just upscaling and MFG, Nvidia has also touted new neural rendering techniques with its Blackwell architecture. In a presentation on neural rendering, Nvidia mentioned Neural Textures, Neural Materials, Neural Volumes, Neural Radiance Fields, and Neural Radiance Cache. It also talked about Mega Geometry in combination with these other rendering features.

Nvidia hasn't announced any specific games that will use these features, instead opting for an Unreal Engine 5-powered demonstration called Zorah that uses the functionality. Using Neural Materials, Zorah was able to deliver higher-quality visuals than traditional materials while using about one-third as much memory. We've seen similar details about NTC (Neural Texture Compression).

The implication here is obvious: an 8GB RTX 50-series GPU could end up behaving like a 24GB GPU. Except, developers will need to explicitly support NTC or Neural Materials; it's not something that Nvidia can enable in its drivers, as far as we can tell—or at least, it's not something Nvidia can enable in its drivers without game developers and publishers opting in for support. (Nvidia doesn't want to upset anyone by doing on-the-fly recompression of game assets.)

That's unfortunate, as universal support for NTC on existing games would go a long way in alleviating concerns with the 8GB and even 12GB Nvidia GPUs. We've seen plenty of games where the RTX 4060 and 4060 Ti 8GB cards fall flat on their faces at higher-quality settings and resolutions. The existence of the 4060 Ti 16GB proves where such issues are due to memory capacity rather than other aspects of a graphics card, and a growing number of games simply don't work well with 8GB cards.

At least for the time being, Neural Rendering won't help—not unless games get patches and maybe even 'Neural Texture Packs' to take advantage of the new features.

© NVIDIA

OTHER GPU UPGRADES

Outside of the core architectural upgrades, Blackwell has some other changes. The video engine now has native decoding and encoding support for 4:2:2 video formats, an increasingly popular choice for higher-end cameras. Processing such content previously would occur on the CPU, and by moving it to the GPU, Nvidia gives an example of a video export going from 110 minutes on the CPU to 10 minutes on the GPU.

Nvidia also upgraded the NVENC encoders for other formats, including the addition of a new AV1 UHQ mode. With the same bitrate, the encoding quality on the 5090 is five percent better than on the 4090 for AV1 encoding, and the AV1 UHQ mode increases that to 10 percent better PSNR (peak signal-to-noise ratio). The RTX 5090 also adds a third NVENC unit (compared to two on the 4090) to boost encoding speed by up to 50 percent.

Nvidia's Broadcast software also has two new modes that are currently Blackwell exclusives: Studio Voice and Virtual Key Light. These aim to improve the quality of voice and face capture, allowing for the use of less than ideal hardware. However, they're also very demanding at present. Virtual Key Light can utilize over 80 percent of an RTX 5080, while Studio Voice needs about 30 percent of the 5080's compute. Additional training and tuning could reduce those requirements, while also improving quality. Our experience with Studio Voice so far has been less than impressive.

Finally, Nvidia enhanced the power gating and power saving features, with faster response times that allow areas of the GPU to enter and exit sleep states much more quickly.

4090, thanks to having 27 percent more compute and 78 percent more bandwidth, and it averaged 25 percent higher performance across an extensive 22-game test suite.

The RTX 5080 was only up to 25 percent faster than the RTX 4080, with an average improvement of just 10 percent—16 percent more compute and 34 percent more bandwidth didn't do all that much. RTX 5070 Ti saw a decent 25 percent average performance improvement over the original RTX 4070 Ti, but only a 13 percent improvement over the 4070 Ti Super. That's to be expected, as it only offers the same amount of theoretical compute as the 4070 Ti Super, with 33 percent more memory bandwidth.

And what about RTX 5070, with a theoretical six percent improvement in compute and 33 percent more memory bandwidth? It was also up to 30 percent faster than the original 4070, but on average only improved performance by 19 percent.

If you look at those figures, it's obvious that the architectural changes to Blackwell are helping, even in native gaming performance evaluations where we're doing our best to make apples to apples comparisons, but gen-on-gen performance only looks to be about 10 percent higher (per SM).

DLSS 4 TRANSFORMERS

The MFG discussion might lead you to reject DLSS 4—where MFG exists—as an unimportant marketing update. However, DLSS 4 contains a lot of additional upgrades that apply to previous-generation RTX GPUs. DLSS upscaling, since version 2.0, has used a convolutional neural network (CNN) to handle upscaling. CNNs delivered some great early results with the modern era of machine learning and AI, but significant improvements have been made with Transformer-based networks. It's the 'T' in ChatGPT: Generative Pre-trained Transformer.

The fundamental idea behind Transformers is that they allow a neural network to focus attention on the most important details. That's glossing over a lot of things, but it's fair to say that Transformers have revolutionized machine learning in the past several years. Now, with DLSS 4, Nvidia is bringing Transformers to its DLSS algorithms.

DLSS 4 upscaling can function in two modes: the legacy CNN-based model, or the new Transformers-based model. That also applies to DLAA, Ray Reconstruction, and even frame generation. The DLSS Transformers model uses four times as much computation as the older CNN model, and delivers superior upscaling, anti-aliasing, and frame generation. How much better can be subjective, but having tested it, the new DLSS Quality mode can be considered better than native in most cases, with Balanced surpassing the CNN Quality rendering. Even the new Performance mode looks very good—not better than native, but probably close to the old Quality setting.

In testing, DLSS Transformers tends to run about 7–10 percent slower than the CNN model on the RTX 50- and 40-series GPUs. Our assumption is that DLSS Transformers will use FP8 precision, and possibly even FP4 for certain calculations on RTX 50-series GPUs. DLSS CNN models used FP16 exclusively, as far as we're aware, but the RTX 30- and 20-series GPUs don't support FP8 in hardware, and the performance hit can be quite a bit larger; around 25–30 percent slower than the CNN model. The same applies to DLAA in supported games.

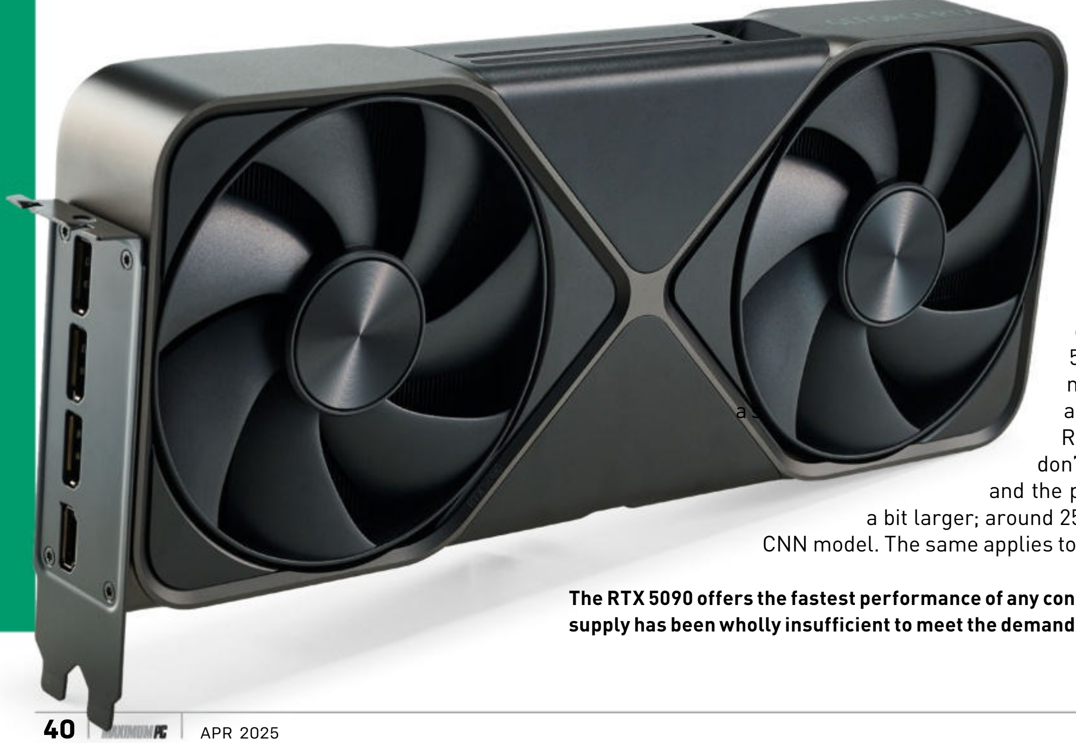

The RTX 5090 offers the fastest performance of any consumer GPU, but so far, the supply has been wholly insufficient to meet the demand, with skyrocketing prices.

The new frame generation model of DLSS 4 is a different can of worms. DLSS 3 framegen used the OFA (Optical Flow Accelerator) fixed function unit to handle frame generation. With Transformers, however, Nvidia has been able to achieve better quality and performance using the tensor cores on RTX 40- and 50-series GPUs. Games with DLSS 4 support will use the new model to deliver better performance even on RTX 40-series cards.

What about running the same models on RTX 30- and 20-series GPUs? It's not supported, and again we wouldn't be surprised if Nvidia is doing FP8 calculations to boost computational throughput to the necessary levels. We also wouldn't be surprised if Nvidia doesn't want to make an RTX 40-series and later exclusive feature work on earlier GPUs. But the bottom line is that framegen on a 40-series GPU should now run about 10–15 percent faster than before, with higher image fidelity as well.

FINAL THOUGHTS

Blackwell RTX 50-series GPUs will power the next couple of years of Nvidia graphics cards. In general, the new architecture promises some noteworthy changes and additions, leveraging AI and machine learning to improve image quality and gaming performance, as well as for generative content for images, text, and videos. Unfortunately for gamers, Nvidia has been too successful in its pursuit of AI over the past several years, and the result is an ongoing shortage of consumer graphics hardware.

The RTX 5090, 5080, and 5070 Ti all sold out immediately at launch, and supplies haven't been remotely sufficient to keep up with demand. We anticipate the RTX 5070 will follow the same path, along with AMD's competing RDNA 4 GPUs that are also arriving soon. The problem is that all these cards, along with Intel's Arc Battlemage, use the same TSMC 5nm-class process node and equipment. There's only so much of that to go around.

Nvidia has sold all its Blackwell B200 data center GPUs though the end of 2025. It can only get about 30 such GPUs per TSMC wafer, but each of those GPUs can sell for tens of thousands of dollars. Faced with a limited number of wafers to go around—with other companies including Apple, Broadcom, Facebook, Google, Microsoft, and more wanting a piece of the action—it's going to be hard for consumer hardware to get significant volumes.

Should Nvidia opt for around $2 million to $3 million of data center products from a wafer, $150,000 of RTX 5090 cards, $145,000 of RTX 5080 cards, or $120,000 of RTX 5070 cards? The math doesn't look good for gamers. Even worse, AMD, Nvidia, and Intel aren't going to waste time building more previous generation hardware either, as it competes for the same limited resources. The RTX 40-series and RX 7000-series GPUs have largely disappeared from retail shelves, and there's not enough new hardware to fill the void.

In such an environment, Blackwell doesn't even need to be amazing to sell out. And fundamentally, we wouldn't classify it as amazing. Nvidia elected to save on costs by not upgrading to a newer TSMC node, but with its data center GPUs accounting for over 88 percent of its $130 billion in revenue from last year, gaming is no longer the top priority. Until and unless the AI bubble bursts, we don't expect things to change.

It's enough to make even the 2020–21 years of GPU Ethereum mining look not so bad. Hopefully we're wrong, and GPU prices and availability will improve. Meanwhile, Blackwell RTX 50-series graphics cards could be the best no one can find in stock. ⏻

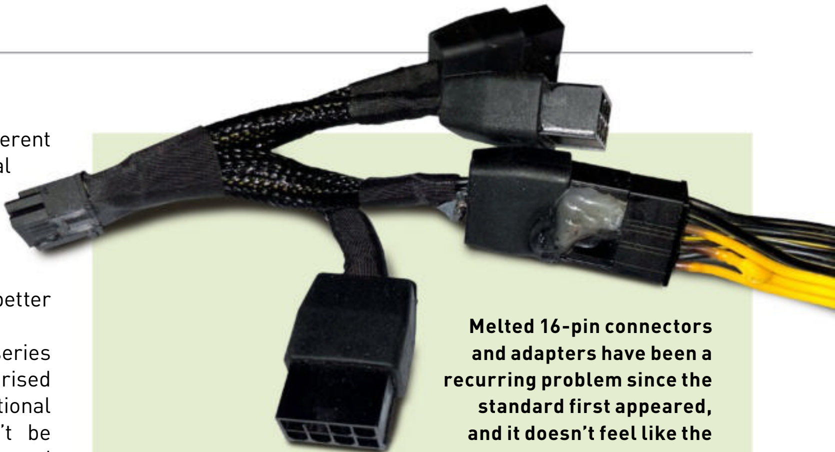

Melted 16-pin connectors and adapters have been a recurring problem since the standard first appeared, and it doesn't feel like the problem is going away.

MELTING 16-PIN, AGAIN?

All the RTX 50-series GPUs released so far require the use of a 16-pin 12V-2x6 connector. The graphics cards all come with the necessary 8-pin to 16-pin adapter, for people that don't have an ATX 3.0 or later power supply with the necessary cable. 12V-2x6 was supposed to fix the problems with melting 16-pin connectors on cards like the RTX 4090, but reports have started to pop up about 5090 cards having meltdowns.

There are some important caveats here, however. First, at least some of the reports come from people using third party 16-pin cables going straight to the PSU. Some of these may still be the older 12VHPWR variant, which was proven to be problematic. Second, the number of reported melting cables—so far—is far lower than what we saw with the RTX 4090. That's despite the RTX 5090 using up to 575W of power for the reference card, with custom AIB models pushing well into the 600W range.

Should you be worried about melting cables with the RTX 50-series GPUs? Our advice is to ensure you have an ATX 3.1 power supply, which means you will have the newer 12V-2x6 16-pin standard on the cables. But those are still hard to find, with only a handful currently shipping. More important than the PSU standard would be to ensure you're using 12V-2x6 cables, and that they are firmly installed.

Frankly, the 16-pin connector feels like a misguided attempt to shrink the power connector size while delivering up to four times as much power. Many electrical engineers have weighed in to express similar thoughts. We probably won't ever see an official recall, but after three years, we're reasonably confident in saying that the 16-pin connector has not proven to be as robust as promised.

The RTX 5090 was up to 40 percent faster than an RTX 4090, thanks to having 27 percent more compute and 78 percent more bandwidth, and it averaged 25 percent higher performance across a 22-game test suite

magazinesdirect.com

Over 100 brands to choose from

GIFT IDEA!
Gifts that keep on giving

3 GREAT REASONS TO SHOP WITH US

- From tech and gaming titles to fashion and celebrity magazines there's something for everyone.
- Immerse yourself in our specialist one-off publications about your favorite hobby or interest.
- Exclusive Offer - Get 15% Off Guides and Specials - **MDMAG15**

VISIT OUR ONLINE STORE AT
www.magazinesdirect.com

✓ Convenient home delivery ✓ Save on shop price ✓ Never miss an issue ✓ 1.4m subscriptions sold

MAXIMUMPC

PREMIUM DIGITAL SUBSCRIPTION

As a premium subscriber you instantly get access to 100+ back issues

GEAR OF THE YEAR
THE 18 BEST PC PRODUC
BUILD
INTEL
First ever Core Ultra build
New Z890 mobo tested
Gaming, AI, and efficiency benchmarked
HISTORY OF RAY TRACING PG.52
AMD vs INTEL GAMING LAPTOPS PG.44
BANISH BLOATWARE FROM WINDOWS PG.60
ARC B580 REVIEWED PG.74
BLUESKY VS X PG.92
BEST MOBILE GPUS PG.12
MAXIMUMPC
MINIMUM BS · FEBRUARY 2025 · www.maximumpc.com
DO IT YOURSELF! STEP-BY-STEP GUIDE PG.16
BUILD AN ALL-INTEL PC
The best-value gaming PC of 2025
Intel's newest CPU & GPU tested
Hyper efficient 1440p gaming
HOW TO USE WINDOWS RECALL PG.60
HISTORY OF AI IN GAMING PG.42
$34.95* for one year/ 13 issues
8TB SSD REVIEWED PG.87
TIMATE
R'S PC
BUILD THIS PC IN JUST 1 HOUR

www.magazinesdirect.com/MPP

*Price in US $. Offer is valid on US orders only, visit us at www.magazinesdirect.com or call 1-844-779-2822 for other subscription options.
Offer valid until April 30 2025.

THE RISE AND FALL OF PACKARD BELL

In this new series, we'll be exploring the companies and tech that made an impact on the computing landscape, but faded away. This month, we'll be covering Packard Bell—a PC brand that dominated desktop computing, and should be familiar to anyone who experienced desktop multimedia in the '90s

THE BEGINNING

The Packard Bell name existed long before home computing ever started. The original Packard Bell Corporation was founded in 1933 by Leon Packard and Herb Bell. It was a producer of radios for homes and cars at a time when radios were covered in wood paneling and cars had swooping curves and chrome.

The company later moved into other consumer electronics, such as television sets, and even military equipment. Of note to this story, in 1957, a subsidiary was formed to produce computers. However, in 1968, Packard Bell was bought out by an American industrial conglomerate, Teledyne Technologies, after which the name was essentially abandoned.

The Packard Bell our readers may know began in 1986. A trio of Israeli-American investors—Beny Alagem, Jason Barzilay, and Alex Sandel—had plans for selling PCs in the US, and bought the rights to the Packard Bell name from Teledyne. Given that Packard Bell had sold computers at one point, it could be argued that the company was legitimately carrying on the brand. But it didn't hurt that the name instantly evoked associations with Hewlett-Packard and Bell Labs.

The old-timey radio brand was perfect for exploiting. It had huge consumer recognition, but few knew what it produced.

The Israeli-American trio's plans were to compete as a mass-production clone maker in the rapidly growing market of IBM PC-Compatibles. Through business connections formed at university, the group was able to establish a supply chain of PC components through Samsung in Korea and Tatung in Taiwan.

Packard Bell released its first PC to market in late 1986. The company's founders used their connections to deliver the best-value PC they could at the lowest price possible. The company also leaned heavily on the acquired Packard Bell brand, capitalizing on consumer nostalgia for instant prestige and brand recognition.

Beyond cheap hardware, they had two main innovations. First, they targeted sales through mass retailers, such as Sears and Walmart, rather than dedicated computer stores. Their other major innovation was to have MS-DOS pre-installed on hard drives, saving customers from having to install it themselves.

These strategies worked. Packard Bell dominated a retail market that established vendors were hesitant to enter, and by 1989, it had annual sales figures of around $600 million. By 1991, it had sales of around $700 million, selling around 500,000 PCs in the US that year, numbers only beaten by IBM and Apple.

THE GOLDEN YEARS

As the mid-'90s approached, the PC world was getting into multimedia. Packard Bell's machines were a mainstay, with funky styling and affordable prices. A shell interface on top of Windows 3.1 was included called Navigator, meant to be simpler to navigate than Windows.

By 1993, Packard Bell had sales of over $1 billion, and had become America's fourth-largest computer vendor. However, other PC manufacturers started adopting its strategy of targeting the low end of the market through standard retail stores.

Packard Bell adapted by adding high-end PCs to its line-up, and began expanding into Europe. New operations were established in the Netherlands, with products selling across 13 countries. An earthquake in January 1994 hit the company's headquarters in California, forcing it to relocate, but Packard Bell seized the opportunity to reconfigure its assembly line process.

By 1994, it had doubled the output of its previous year, and 1995 saw Packard Bell become the biggest PC seller in America, with 13 percent of overall US sales.

© WIKIPEDIA

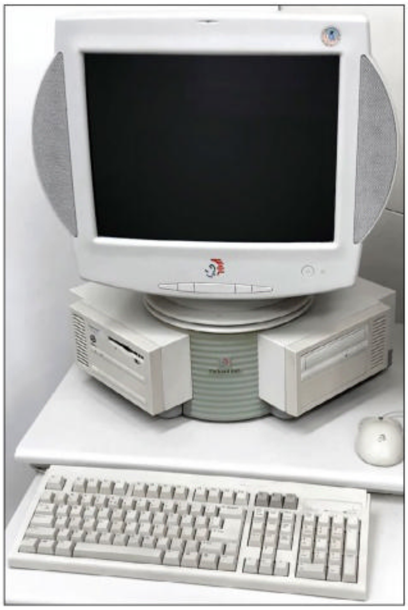

Packard Bell's multimedia PCs had wacky form factors, like this Corner Computer.

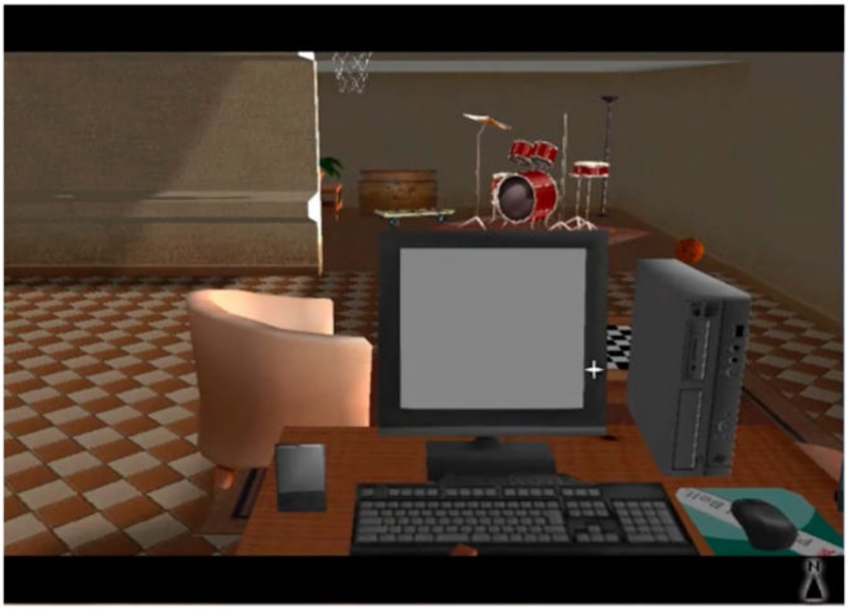

Packard Bell's Navigator interface eventually got the '90s interactive 3D treatment. If anyone genuinely used it daily, send us a postcard!

This would, however, be the peak of its success. Despite phenomenal sales, the company had been relying on thin margins and several missteps meant the money was about to run out.

THE CRACKS START TO APPEAR

In April 1995, a legal battle broke out between Compaq and Packard Bell. Compaq accused Packard Bell of installing used parts in its PCs and selling them as if they were new. Packard Bell *was* using refurbished parts but insisted that it was conforming with standard industry practices by doing so. Company lawyer, Marshall Grossman, stated:

"Computer chips have a lifetime of 1,000 years. For a company to throw such chips into the trash would be an obvious waste of money."

In September 1996, it was ruled that Packard Bell had to pay out more than $5 million in multi-state and federal settlements. It was required to label its PCs, stating that they may include used parts, and to warranty them as if new.

Beyond the damage of the used parts scandal, Packard Bell had also been making heavy investments, such as the acquisition of US company Zenith Data Systems, and taken a big hit when it miscalculated which Pentium chips to order in preparation for Christmas.

Smelling opportunity, Japanese electronic giant, NEC, bought out Packard Bell in June 1996. The combined firm was named Packard Bell NEC—an $8 billion operation that could unlock different markets for both firms.

NEC had a strong corporate presence in Japan and its PC-98 line was highly revered. However, its prestigious computing line was also a proprietary Japanese technology—NEC needed a strong entry into the international PC market with the now ubiquitous standard of Windows on x86.

At the time, despite falling revenues, Packard Bell was second only to Compaq, with sales somewhere between 3.2 to 3.4 million a year. It could provide an instant worldwide network for PCs and economies of scale, while its purchase of Zenith Data Systems promised strong sales in the government market.

THE FALL

Initially, the partnership appeared to work well, with the conglomerate selling almost one in seven PCs in the US market. NEC was willing to tolerate some operating losses, but by late 1998, things looked grim. The industry was flooded by low-cost machines sold directly to customers, bypassing normal retail channels.

WHAT LAPTOP
Packard Bell

During the '90s, Packard Bell dominated desktop multimedia PCs. In the 21st century, the marque has mostly sold budget laptops outside North America.

The Packard Bell NEC share of the US market had dropped from 14.1 percent to 5.5 percent. The partnership lost $650 million over the year, and more than $1 billion over 1997 and 1998.

In November 1999, NEC Japan announced that it would lay off most of Packard Bell's staff—from 2,600 to just a few hundred—shutter its plant in Sacramento, and retire the brand in the US. All but one of the senior executives resigned.

Years later, in January 2008, 75 percent of Packard Bell's parent company, PB Holdings, was bought by Taiwanese manufacturer, Acer, for $48.5 million. By now, Packard Bell had truly lost its American image, with a CNET article calling the company "European PC vendor Packard Bell" and a Network World article referring to it as "Dutch." Acer's focus for Packard Bell would be markets in Africa, Europe, and the Middle East.

In 2018, the brand was relaunched—to little fanfare—in North America by Brooklyn-based electronics manufacturer Southern Telecom. They currently sell Packard Bell-branded monitors, laptops, and tablets.

The Packard Bell brand is still in use by Acer for the South African market, where it sells a range of budget laptops. ⏻

© WIKIPEDIA, WWW.TECHRADAR.COM

ITX
{A HISTORY}

FOR THOSE
ROG
STRIX
WHO DARE
ROG
STRIX
FOR THOSE WHO DARE

From conception to evolution, how ITX changed computing

ITX motherboards are amazing, when you think about it. There's a similarity between them and that age-old technology adage: things get smaller, and more powerful as they do so. It's that Moore's Law-esque belief that holds precedent in the very idea of ITX. From cellphones to laptops to supercomputers and more, there's a similar vein that runs through it. That's arguably why the ITX and small form-factor fan base is as passionate as they are.

ITX breaks the trend. It's different; the very idea that you can take super-small systems, build a PC in a suitcase, or have a chassis the size of a toaster, and effectively cram it with the same top-tier hardware you'd find in a gargantuan full-tower and have it remain portable, while also being well-cooled, is awesome. It's a sleeper build in disguise.

It comes with its own challenges, of course; ITX builds are difficult to work with, build in, and often have major drawbacks in terms of connectivity and I/O (you can read our build feature this month to find us experiencing exactly that). And yet, there's a charm to it all.

So where did it all begin? We know the world of form factors—certainly for motherboards—is a mess. There's no consistency, of course; what actually is a form factor and what isn't is wildly contested, but for ITX, things are a little more clean-cut than you might think. Let's dive in, shall we? –ZAC STOREY

The arrival of the form factors

Getting from here to there

The world of the humble motherboard is one of weird, wacky, and wonderful drama. It's shadowed in myth and legend, and submersed in unexplained Lovecraftian rituals. Well, not quite, but compared to every other form factor, standard, and more, it's very much a 'flying by the seat of your pants' affair.

We've already covered the origins of these form factors fairly recently, but in brief, IBM was at the forefront of mobo form factors from 1983 all the way up to 1995. It introduced XT (8.5 x 11 inches), AT (12 x 11-13 inches), and Baby-AT (8.5 x 10-13 inches). In 1995, Intel took the reins, and began to introduce the more common form factors we know and love today. First came ATX in 1995 with its 12 x 9.6 inches of glory, then Micro-ATX in 1996—a slightly shorter 9.6 x 9.6 inches.

THE ARRIVAL OF COMPROMISE

Although Intel continued to launch form factors with both BTX and FlexATX, it wasn't actually the company that introduced Mini-ITX onto the world stage, despite ATX and Micro-ATX continuing to have the dominance that they do. No. VIA Technologies debuted Mini-ITX to the world in March of 2001.

VIA Technologies was originally founded in California back in 1987 by Cher Wang, before the company relocated to Taipei in 1992. It specialized predominantly in the manufacturing of chipsets and processors. VIA Labs, a subsidiary company, generally develops controllers for USB platforms and peripherals as well, and still employs well over 2,000 people from across the planet.

Left: Asus's X870-I is a fine example of how far we've come on ITX design. It's all yours for the asking price of $520.

Raised PCB, anyone?

As for ITX, that's the brainchild of Robert Kuo, current Assistant Vice President at VIA Technologies. Robert was chief R&D expert at the company, and introduced the 215x191mm ITX reference board to much chagrin from the competition. It wouldn't be until Intel took the platform on, introducing an array of ITX motherboards in 2008 with its Atom line of processors, that things really started to kick off. At IDF in Shanghai, Intel launched its first generation of Atom processors, all of which were aimed at low-power entry-level PC systems for schools and other applications. However, these were soldered; the first board was the D945GCLF, running a single-core Atom 230 @ 1.6 GHz and one slot of DDR2.

We've done as much digging as we can, and the first ITX boards we can find with an actual socket of any note, rather than a soldered chip, are Zotac's GeForce 9300 ITX WiFi motherboard and the 8200-ITX WiFi Board. There are a few other boards from now-deceased manufacturers, but the Zotac ones in particular made the biggest dent on the market.

These were specifically designed with the home theater PC in mind for those who didn't want to sacrifice on CPU performance directly. The 9300 was an Intel offering, while the 8200 was specifically for AMD. There's not a huge amount of detail out there on them. Both featured soldered GeForce Nvidia chipsets on there to drive home theater entertainment, and both had wireless integration as standard as well, with support for two DDR2 slots, as you'd

RECAPPING MOTHERBOARD FORM FACTORS

Year Introduced	Acronym	Expanded	Origin	Dimensions (inches)
1995	ATX	Advanced Technology Extended	Intel	12 x 9.6
1996	Micro-ATX	Micro Advanced Technology Extended	Intel	9.6 x 9.6
2001	Mini-ITX	Mini Information Technology Extended	VIA Technologies Inc	6.7 x 6.7
2003	Nano-ITX	Nano Information Technology Extended	VIA Technologies Inc	4.7 x 4.7
2009 - 2010	XL-ATX	Extra Large Advanced Technology Extended	EVGA, Gigabyte, MSI	13.5-13.6 x 10.3-10.4
???	E-ATX	Extended Advanced Technology Extended	Various	12 x 10-13

© ANANDTECH

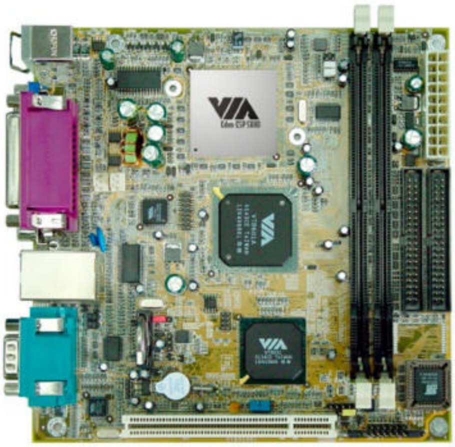

The first ever ITX motherboard produced by Via Technologies.

Blast from the past, Intel's first Atom ITX boa‹rd produced in 2008.

Zotac's GeForce 8200-ITX WiFi was an epic ITX option.

probably expect. The 8200 was an AM2+ socket, and the 9300 an LGA 775. Each had SATA/300 ports, eSATA, and both also featured a PCIe Gen 1 x16 slot port as well, allowing for the introduction of dedicated graphics cards of the era.

They certainly weren't classed as premium motherboard solutions at the time, despite their compact form factor. Over time, the form factor would develop to become the juggernaut that it is today, with Intel's Z77 chipset in particular bringing an array of new ITX boards to the forefront. Our personal favorite was ASRock's Z77E-ITX (the board this journalist overclocked his Core i5-2500K to 5.2 GHz on).

ASRock continued pushing the boundaries of what was possible as well, and in 2015, ASRock debuted the first ever enthusiast-grade motherboard, the X99E-ITX/AC. It was complete with a six-phase design, 802.11ac WiFi, USB 3.1 Gen 2 (I actually reviewed it in my first ever feature back in the September 2015 issue), along with support for DDR4 at 3200 MT/s and, of course, a PCIe 3.0 x16 slot as well. Theoretically, back in the day, you could chuck one of Intel's Haswell-E Extreme Edition chips in here, the Core i7-5960X, which had a staggering eight cores, 16 threads, 3.5GHz clock speed, and 20MB of L3 cache with a wild 140W TDP, retailing at $1,000. How times have changed. In fact, the only major downsides to the board were the odd cooler mounting locations, which weren't in the traditional positions for an LGA2011 chip, and the fact that it only supported dual-channel memory, where one of the big features of X99 was its quad-channel support.

That latter fact was something ASRock would rectify with its X299E variant instead. With the X299E-ITX/ac, the new LGA2066 socket now supported that quad-channel memory design, but under one condition: you had to populate it with SODIMM memory instead. It also featured some of the very first M.2 PCIe 3.0 SSD slots on a motherboard, with two on the back, along with the addition of a vertical PCB to allow for an additional M.2-2280, plus some extra connectivity and additional hardware. Rear I/O suffered as a result, and it only featured five USB 3.1 Gen2 Type A ports and one Type C port. AC WiFi was included, and it did come with dual Gigabit Ethernet, but that was it.

A massive leap over its predecessor, you could theoretically kit out the X299E-ITX with Intel's Core i9-7980XE, complete with its 18 cores and 36 threads, plus a 4.4 GHz clock speed, and 44 PCIe lanes—admirable even by today's standards. It wasn't cheap, though, and it sauntered in with an RRP of $2,000 alone (for the chip).

Since then, each subsequent generation of chip release has had a host

CHASSIS DESIGN OVER TIME

As the motherboard has developed over the last decade, chassis design has followed suit. In the beginning, most ITX cases were designed with that ethos in mind. Either they were all-in-one solutions on the back of a monitor, or home theater PC boxes no bigger than a KFC bucket. That inspired a lot of early system builders to experiment with ITX, building all manner of PCs into suitcases, soccer balls, toasters, older cases, or anything they could find.

Case manufacturers, however, were more advanced. BitFenix in particular introduced its Prodigy line, complete with a 230mm front intake fan, providing an ungodly amount of cool airflow to the internal hardware inside. Corsair similarly had its radical 380T with support for full-sized 240mm liquid-cooler as well. In fact, most of the early year cases were designed with LAN parties in mind.

It was only in the late 2010s that chassis design got more serious when it came to the ITX form factor. With the launch of the Corsair One, and then NZXT's DIY H1 line, ITX cases became more luxurious, bigger, and bolder, with support for 360mm coolers from some manufacturers. Cooling became a dominant design element, particularly as the CPU war between AMD and Intel heated up.

As a result, there are ITX chassis for every use case imaginable. From home theaters to liquid-cooled monsters, LAN parties to gaming machines, and console replacements to professional productivity rigs, rendering away late into the night.

© CORSAIR, VIA TECHNOLOGIES

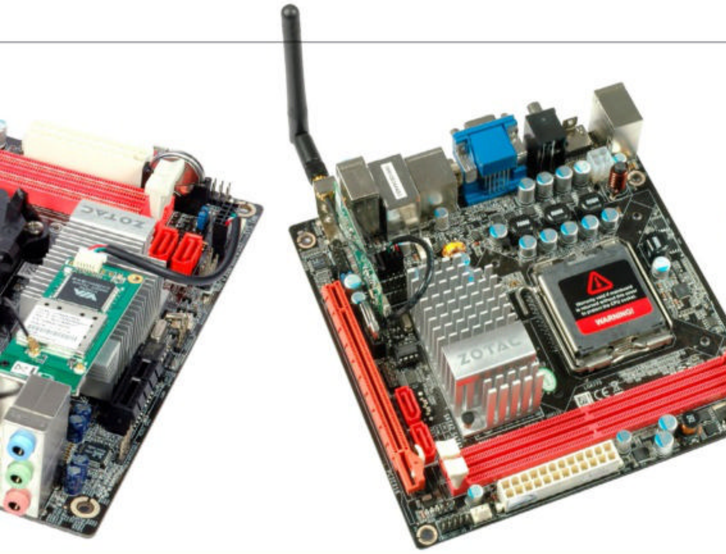

Zotac isn't a name associated with motherboards, but here's its 9300.

ASRock was the first to push the boat out and develop this insanity.

of ITX motherboards to back it up in ever-increasing numbers, and as the form factor continues to grow in popularity, so too has the number of cases available supporting it.

ASUS AND THE RAISED PCB

ITX's biggest problem was the design ethos behind it, mostly stipulated around quiet, low-power operation systems that weren't kitted out or designed for the high-powered gaming and rendering world. Manufacturers knew this, yet the demand for high-performance ITX motherboards was quite high. This biggest issue that stymied these attempts was the lack of good power delivery. The problem, however, was that the motherboards didn't have the capacity to feature a fully fledged VRM solution like their ATX counterparts could muster. There just wasn't enough space on the board.

That led to Asus's Maximus VI Impact. The company's first Impact motherboard, based on the Z87 chipset, featured a raised PCB, complete with an 8+2 power phase design, utilizing 60A choke, along with all the features you'd expect from a gaming-oriented motherboard of the time. This effectively ensured your LGA1150 CPU could operate at its turbo boost frequencies for longer without throttling, and allowed for some fairly healthy overclocking as well.

THE END?

ITX has changed a lot over the years, becoming incredibly advanced compared to what it first set out to be. This low-powered, super-quiet budget solution, designed with educational and basic office needs in mind, designed by a relatively small company in Taipei, has evolved into a glorious catalyst, capable of allowing any would-be PC enthusiast to downsize and pack as much punch and power into their PC as they could possibly provide.

Since the advent of Z77 and onward, gamers can cram top-end processors, the fastest storage, and the best graphics cards into any small form factor chassis they can find. With it, you can effectively drive performance that put full-tower behemoths to shame.

Over the last decade, *Maximum PC* has built all manner of weird and wonderful machines, but without a doubt, it's the ITX numbers that time and time again have proven to be the most intricate, odd, beautiful, and powerful ones of the bunch. The Dream Machines might hold the spotlight, but it's the ITX builds that we remember most. There's something magical about these things, and no matter how this industry progresses, knowing that ITX, and custom PC building will still have that scrappy fiery step-child ripping up the rulebook and holding power to account, fills us with joy. ⏻

WHAT DO YOU NEED FOR ITX?

Honestly, not a lot. The form factor is probably one of the most forgiving out there. That's arguably a reason it has stayed around for as long as it has. In short, you only need to account for three main things. First, the dimensions must be 17 x 17cm (or 6.9 x 6.9 inches); the mounting hole locations must be in the same position as the four found on the top left side of an ATX motherboard, and lastly, the rear I/O must be in the same position as well. Everything else is fair game—socket position, DIMM slots, headers, power, you name it. Via technologies were incredibly generous in that way, and it has certainly given ITX longer legs than the likes of FlexATX, Mini ATX, and similar form factors that emerged at the time.

MAXIMUMPC
WATER-COOLED OVERCLOCKING
MAXIMUMPC
OPTIMIZE WINDOWS
MAXIMUMPC
FORGE THE PERFECT EDITING PC
MAXIMUMPC
SUPER NANO COOLING
MAXIMUMPC
MICRO BUILD, MAX POWER!

ITX has been a recurring feature on our covers over the years.

© ANANDTECH

Corsair Vengeance A7500

AMD'S RYZEN 7 9800X3D CPU isn't just the gaming CPU du jour; it's becoming a long-term legend. It's no surprise, then, to find it at the beating heart of Corsair's Vengeance A7500 desktop rig. As a cutting-edge gaming PC, it's also no shock to find Nvidia's latest RTX GPUs at its heart. These days, buying a prebuilt system like this might be one of the only ways to get a decent graphics card.

It's actually the rest of this rig—as opposed to the industry-standard, albeit excellent, core components—that really stand out. Corsair is unusual in being a system builder that also manufactures a wide range of PC components itself. Most PC builders might commission their own case design, then cobble together a bunch of third-party parts. Not Corsair; it has far better vertical control of a desktop PC's engineering.

Not only do you get Corsair's 3500X Mid-Tower Case with a wraparound glass aesthetic; the Vengeance A7500 is also powered by a Corsair 850W 80+ Gold Fully Modular power supply, runs 32GB of Corsair Vengeance DDR5 6400 RAM, and sports a 2TB Corsair PCIe 4.0 SSD. Nice.

There's also the Corsair Nautilus 240mm liquid CPU cooler. It's a low-noise approach to highly efficient CPU cooling, thanks to RS120 fans that deliver high airflow, static pressure, and low noise, courtesy of precise PWM fan speed control with Zero RPM mode.

Currently, you can choose from any of Nvidia's new RTX 50 series GPUs. We'd say the RTX 5080 option for $3,499 all in looks like a great choice, especially when you consider that the step up to RTX 5090 graphics forces a doubling of RAM to 64GB and a price of $5,999. Ouch! –MAXIMUM PC

1 NVIDIA FOR THE WIN

You can choose from any of Nvidia's latest RTX 50 series graphics. Indeed, with graphics cards currently in such short supply, rigs like this are just about the only way to get an RTX 50 GPU right now. Our pick is the RTX 5080. Going for an RTX 5090 forces a few other upgrades and results in a painful price premium of around $2,500.

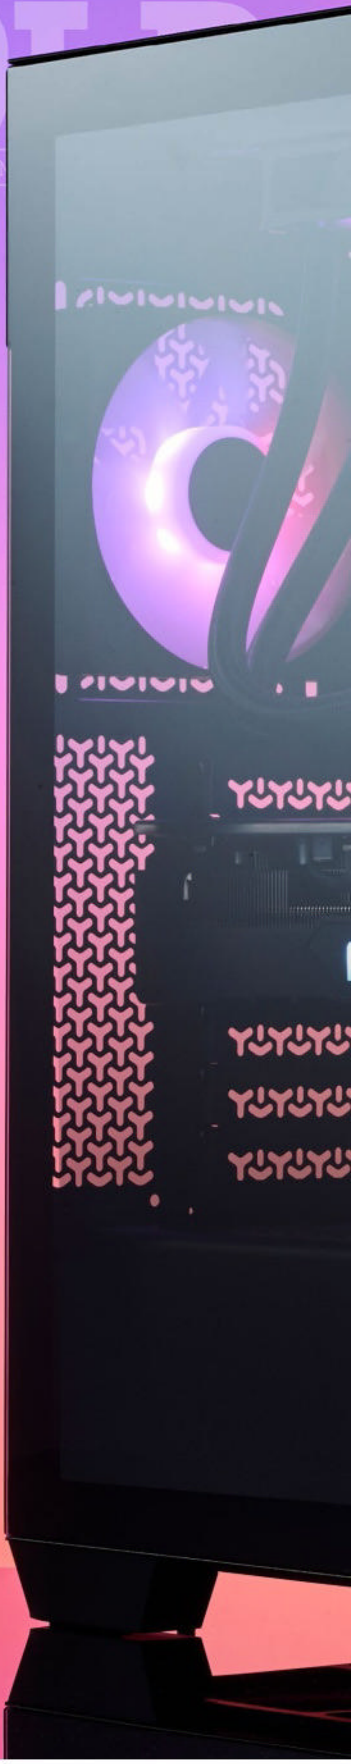

VENGEANCE
msi
PNY
GEFORCE RTX

2 CORSAIR CORSAIR

Unlike most PC builders, who cobble together systems using third-party bits, Corsair makes most of the critical PC components outside of the CPU and graphics card itself. That means you're getting a nearly pure Corsair rig, including RAM, SSD, cooling, PSU, and the case. So, there's much more engineering integration than you get from most of the competition.

3 COOL TRICKS

Among the Corsair-engineered highlights is the Corsair Nautilus 240mm liquid CPU cooler. Apart from lots of low noise features, it also sports a convex CPU block. The idea is to ensure optimal contact with your CPU's integrated heat spreader. It also comes with paste applied in an optimized pattern for easy installation and excellent thermal performance.

BUILD A STEAM DECK

John Knight tries one of several SteamOS imitations, and is overwhelmed by the future possibilities for Linux PC gaming

THE LINUX-BASED Steam Deck has revolutionized PC gaming with a console-like experience in handheld form. Valve has been promising to release SteamOS for general PCs—beyond just the Deck—for ages now, but some folks have got tired of waiting, and decided to take matters into their own hands.

Whether you can't afford a Steam Deck or just want SteamOS on another device, maybe one of these SteamOS imitators is just what you're looking for.

There are a number of competing gaming distros, several of which are SteamOS imitations. Right now, your choices are HoloISO, ChimeraOS, and Bazzite. All three are worthy of your attention, but we have chosen Bazzite because it supports the broadest range of hardware and form factors.

Unlike HoloISO and ChimeraOS, which attempt to replicate the Steam Deck's Arch-based underpinnings as closely as possible, Bazzite is based on Universal Blue, a set of OS images based on Fedora Atomic, itself a respin of Fedora, centred around an almost-immutable atomic system. A word of warning: it's still early

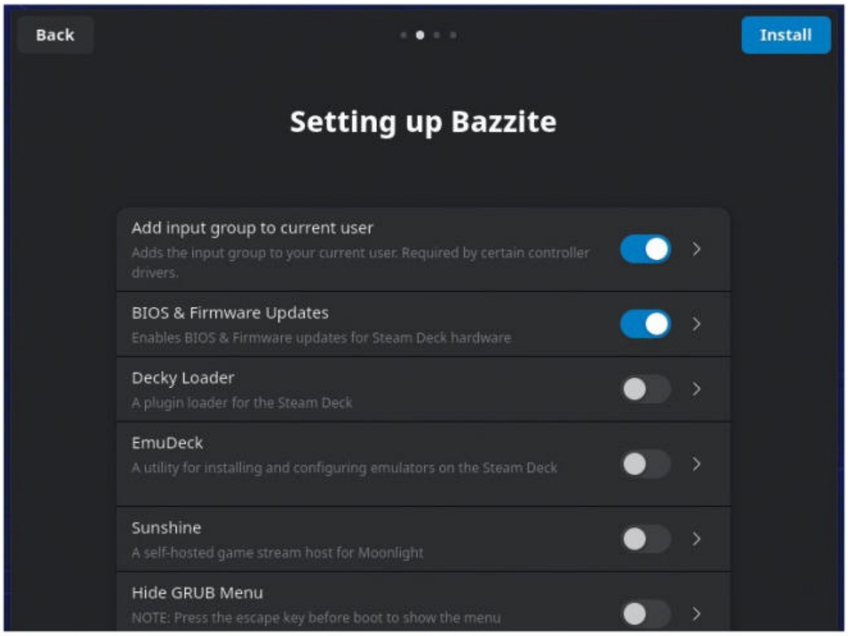

Setting up Bazzite

- Add input group to current user ●
  Adds the input group to your current user. Required by certain controller drivers.
- BIOS & Firmware Updates ●
  Enables BIOS & Firmware updates for Steam Deck hardware
- Decky Loader ○
  A plugin loader for the Steam Deck
- EmuDeck ○
  A utility for installing and configuring emulators on the Steam Deck
- Sunshine ○
  A self-hosted game stream host for Moonlight
- Hide GRUB Menu ○
  NOTE: Press the escape key before boot to show the menu

With a broader choice of hardware, pre-installed tweaks, and a choice of KDE or Gnome, Bazzite may be more appealing than official SteamOS.

days for these distros, and Bazzite is still a bit of a buggy beta, but it's worth persevering nonetheless!

Bazzite (https://bazzite.gg) has minimum requirements for its Deck-like gaming mode, but provides separate desktop-only images for legacy hardware. We will be focusing on the full-fat gaming version in this article, but will touch on the legacy desktop edition later.

According to the documentation, the main prerequisite is 'a graphics card that can utilize Vulkan 1.3+.' For AMD GPUs, that means a Radeon RX 400 (2016) series or newer. For Nvidia users, it's an RTX Series or GTX 16x0 (2018). Intel graphics users need an Arc GPU, otherwise you will be stuck with desktop-only mode.

To get started, head to the Bazzite website, and click the Download Bazzite button in the top-right corner. Look for What Hardware Are You Using? and click Select Your Hardware. This gives a list of different form factors, including laptops, handhelds, and desktop PCs—we went with the home theatre PC option.

This activates two new drop-down menus: Who Is The Vendor Of Your Primary GPU? and What Desktop Environment Do You Prefer? We've already covered the GPU issue, but if you pick one of the current options, another menu appears asking whether you want Steam Gaming Mode. Choose Yes unless you are stuck with old hardware. Select the desktop environment you want—KDE is what the Steam Deck uses, but Bazzite also provides Gnome, with Budgie and Cosmic coming later. Once you have made all your choices, a big button appears with a download link, providing an ISO file.

After the ISO has downloaded, you need to flash the file to a USB flash drive. Plenty of flashing tools are available, but a popular choice is balenaEtcher (https://etcher.balena.io), which runs on Linux, Windows, and Mac OS. After the flashing has finished, you can close the program and restart the computer, ready to boot from the flash drive.

When the computer restarts, you need to enter the BIOS to configure some boot options and boot from the flash drive. To enter the BIOS, you usually have to hold down either Delete or F2 when your computer starts. Once inside, enter Advanced Mode to change your boot options, under which there may be a Boot section—this varies depending on manufacturers.

The documentation says to disable CSM (Compatability Support Module), which forces UEFI mode. You also need to disable Secure Boot. This may have a different name in your BIOS, but under Secure Boot, change OS Type from Windows UEFI Mode to Other OS.

If you have managed to find and set these options, the surest way to proceed is to save the changes and exit the BIOS (this restarts the computer). Then go back into the BIOS and choose your flash drive as the first bootable device, save the changes, and exit again. You may also be able to just click on the flash drive in a Boot Override menu.

GATHER SOME STEAM

Once the flash drive boots, a GRUB menu appears on screen bearing the text Install Bazzite-Deck 40 (or whatever version number exists now). It then goes through all the usual boot stuff before starting a graphical installer. The installer is taken from Fedora, which is okay overall, but still feels like something from 20 years ago.

Starting with Installation Destination, you have to choose a hard drive from a list of Local Standard Disks. When you select the drive, you can choose between Automatic and Custom partitioning. Click Done to proceed.

If your disk is pre-used, Automatic will probably complain about not having enough room, but then give you the option to Reclaim Space. This opens a window that allows you to delete or shrink partitions. We chose Delete All, but you may try shrinking partitions if you want to dual boot. Click Reclaim Space in the bottom-right corner of the window when you're done.

LEGACY DESKTOP VERSION

We tried the desktop-only version of Bazzite for old hardware, using only drivers that came with the distro. Using some random parts, our test system had an Nvidia GeForce GTX 1050, 16GB RAM, and an Intel Core i5-3470 @ 3.2GHz. [Is that *really* legacy?—ed]

Unlike the standard edition of Bazzite, the desktop-only version boots to a standard KDE or Gnome login screen. Once logged in, Steam starts updating automatically, which happens behind the same first-run wizard as Bazzite's normal mode.

KDE and Gnome were smooth and responsive, and passed the screen-tearing test on YouTube, despite us not having installed any special drivers. Additional hybrid Nvidia drivers were added to the first-run wizard process: Supergfxctl (Hybrid GPU Switching). Big Picture Mode was very slow and clunky, despite a 3D title like *Alien: Isolation* running quite well. An older, slower AMD Radeon HD 7800 was faster and more stable, with no glitching.

© HTTPS://BAZZITE.GG

A cheap mini PC can be turned into a surprisingly versatile console for sitting room gaming.

The Create User section has an old-school dialog for setting up your user account and password—otherwise, it makes a generic 'bazzite' user account. Enter your name, username, and password, then choose if you want to add administrative privileges (leave this checked if you don't know what you're doing), and check or uncheck the box on whether you want to require a password to use this account (again, check this unless you have a reason not to).

To connect to Wi-Fi during installation, open Network & Host Name. You will find your network devices under Ethernet or Wireless. Choose your Wi-Fi adaptor, and click Select Network to configure it.

When you're finished setting up, click Begin Installation, and the progress bar starts as it copies the system to the hard drive. When the process is finished, click Reboot System. You might encounter a screen that says 'Perform MOK management' when you reboot. Just choose Continue Boot if so.

FIRST RUN SETTINGS

GRUB boots again. If there are two menu entries, just choose the first, or wait, as there are various updates. Then there is a welcome screen, followed by a choice of languages. This is followed by yet another update screen that may take a good 15 minutes or so, depending on your processor and connection speed. After this, we got the Choose Network dialog again, then a screen saying the system was resetting.

Finally, after this reboot, you see a Steam login screen. Here, you enter your Steam account details to continue. You probably have to enter an authentication code that is sent to your email address—welcome to the faff of modern computing. Now, at last, there is a screen that says 'Welcome to Steam Deck'.

ENTER GAMING MODE

If you've never used a Steam Deck menu, you can open it with the Steam button (the big button in the middle of your controller).

To install some games, open the Steam menu, and select Library or Store. Whether browsing the store page or your own library, the default category opened is Great On Deck, which guarantees a better playing experience. Otherwise, you can browse the rest of the collection in Store Home, or All Games if you're looking at your own library.

To install a game, either left-click it or press the A button on your controller. This opens the game's entry page. If you're using a controller, move the stick around to jump between menu entries, then press the blue Install button.

For non-Linux games, a Steam Deck Compatibility box appears. If it's not on the recommended game list, an info box appears with the game's current support status, saying what is known to work and what features may still be unsupported, like integrated Steam inputs or text scaling. Click Install if you wish to proceed—you can continue browsing while a game is downloading.

Once the download has finished, at the top of the screen is the Installed tab, where you can open the game and click the green Play button.

IT'S A DESKTOP!

To switch from the default Gaming Mode to Desktop Mode, open the menu, choose Power, and from the menu, select Switch To Desktop. In imitation of the Steam Deck, a clean and minimalist KDE desktop is provided. A first-run wizard provides a bunch of cool tweaks to get you gaming faster. More on that later.

Browser-wise, a stable version of Firefox is included. YouTube worked immediately, without screen tearing. The all-powerful game launcher and Wine curator, Lutris, is obviously included ,too.

If you take your time to go through the first-run wizard, other software is ready to download. Browsers such as Chrome,

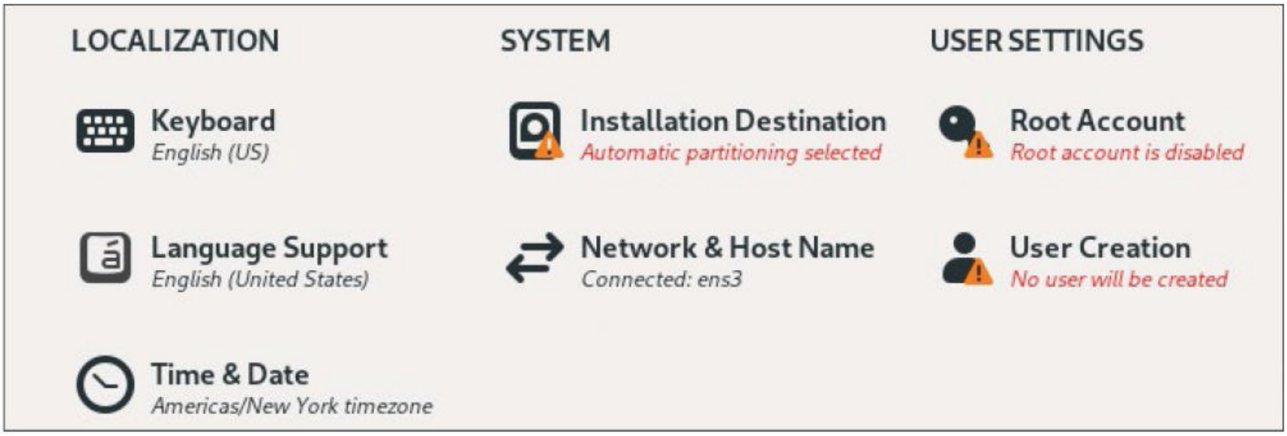

LOCALIZATION

Keyboard
English (US)

Language Support
English (United States)

Time & Date
Americas/New York timezone

SYSTEM

Installation Destination
Automatic partitioning selected

Network & Host Name
Connected: ens3

USER SETTINGS

Root Account
Root account is disabled

User Creation
No user will be created

Bazzite's antiquated Fedora installer can be a bit intimidating, but with a little hand-holding, you'll be okay.

© HTTPS://BAZZITE.GG

Brave, Vivaldi, and Opera are ready to go, as are emulators, multimedia, office apps, and more.

If you're using a desktop-only edition of Bazzite, and wish to emulate the full-screen feel of gaming mode, click View > Big Picture Mode in the Steam client.

A big icon is on the desktop—Return To Gaming Mode—which also has a warning prompt, in case you click on the icon by mistake.

MAKING IT HEROIC

There is far more on offer if you're willing to switch to Desktop Mode and explore, mostly available via the first-run wizard.

Of particular note is the Heroic Games Launcher, which allows you to run games from GOG, Epic, and Amazon. We were especially impressed by Heroic's work on GOG, as old DOS and Windows games could be installed and run with very little effort. These were then integrated into our Steam library as though they were normal Steam titles, and we were even able to choose between using Wine or Proton to launch them.

Some other cool features are an Itch store client, Minecraft builds, a GeForce Now client, Oversteer, a steering wheel manager, and Waydroid, for those who want to use Android apps.

Any retro gamers should explore EmuDeck. An extensive ROM collection may take a while to set up, but it can install an impressive number of emulators and integrate emulated games into your Steam library. Much like Heroic's work with GOG, we were able to integrate Sega Mega Drive games into our library and launch them like regular Steam titles.

Serious gamers should try ProtonUp-Qt, which lets you download custom versions of Wine or Proton, and then integrate them into Lutris or Steam.

However, the most transformative tool is Decky Loader. This has its own store of plugins that bring more functionality and customization to SteamOS—Decky Loader could warrant an article of its own, and is a must-have for any Deck owner.

TACKLING ISSUES

You might want to keep a keyboard and mouse plugged in for the first hour. We were using an old Xbox 360 controller for testing, which worked when gaming, but only partially worked with the Steam menu interface. The menus all worked fine after we launched our first game, but that first run might be buggy.

When we logged out of KDE in the regular edition's desktop mode, we got a black screen and blinking cursor—the Bazzite team may never have tried this!

© PLAYDEAD, DOUBLE ELEVEN

BALATRO
PLAY
OPTIONS
QUIT
COLLECTION
ROG

According to YouTuber TechDweeb, once it has been improved with Bazzite, the more powerful ROG Ally X Windows handheld feels like a Steam Deck 2.

Likewise, if you have anything slightly complicated in the sound department, you may not have sound out of the box. Our Asus soundcard required entering Desktop Mode and changing outputs in KDE's control center.

Despite initial teething problems, once we had the system up and running, the experience was eye-opening. SteamOS may make sense on a handheld, but it makes even more as a sitting room gaming system. Being able to play your PC games on a sofa with a controller has appeal, but it's even better when you can switch between gaming and desktop modes, and they just get out of each other's way.

If this kind of setup (whether SteamOS proper or a rival like Bazzite) becomes a regular out-of-the-box retail experience, desktop gaming is going to become more niche, and traditional game consoles are going to have their work cut out, as this kind of system will be very difficult to argue against.

2D platformers like *Gris* and *Limbo* are modern classics worth just as much attention as the next triple-A blockbuster.

LOW-END HARDWARE

You don't need a powerful computer to have a good gaming experience. For starters, the PC has the biggest library of any platform, and any full-time Linux user probably has an enormous backlog of games to catch up on from the last two decades.

Beyond older games, a slower PC may be suited to the modern indie scene, which is in something of a golden era right now, and a great alternative to today's soulless triple-A gaming selection. Modern 2D platformers in particular are often criminally underrated, bursting with originality, and run on just about anything.

If nothing else, a low-powered box in the sitting room could stream triple-A titles from a more powerful work PC. Sometimes, gaming is just better in elastic-waisted trousers, lying on the sofa with a packet of Quavers.

Bag a software bargain

Don't pay full price for software when we can offer you huge reductions on everything from security software to Windows 11 Pro Upgrade

Norton 360 Premium 2025 for $29.99

WE'VE NEGOTIATED a killer deal with Norton. No subscriptions; just a one-off bargain price of $29.99, compared to the regular $109 charge. That buys you 15 months of cover from the Norton Security suite across 10 devices

Because it's the Premium version, you get a host of extra tools. There's 75GB of cloud storage for secure backups, plus the full version of Norton Secure VPN, both to protect your identity and provide a way of watching TV from home while abroad.

Norton Parental Control provides tools to see exactly what your kids are up to on their various devices, and you also get GPS location monitoring for Android and iOS.

To round things off, Norton Password Manager generates and stores passwords across all your devices, while SafeCam for PC stops cybercriminals taking photos on your webcam without your knowledge.

15 MONTH PRE-PAID SUBSCRIPTION
norton
10 DEVICES
Norton 360 Premium

SAVE 72%

ALSO CONSIDER **Avast Ultimate Suite 2025 for $29.99 (10 devices for two years)** **SAVE 87%**

Windows 11 Pro Upgrade for $59.00

WANT TO UPGRADE from Windows 11 Home to Windows 11 Pro? This is a simple upgrade from Home to Professional that will then offer you the additional Professional features for a competitive price. What's more, the upgrade occurs as soon as you restart your computer.

Pro includes business functionality for encryption, remote login, virtual machines, and more. Microsoft charges a staggering $219.99 for Windows 11 Professional. By choosing an OEM license—and avoiding the many dubious sources of licenses you may find online—you cut that price substantially.

Windows 11
Pro Upgrade

EXTRA 20% OFF!

USE **PROMAR20** AT CHECKOUT

ALSO CONSIDER **Windows 11 Home for $49.99** **SAVE 64%**

CCleaner 6 Professional $14.95

CCLEANER 6 PROFESSIONAL ensures that your PC keeps running as smoothly as it did the day you bought it. This tool will remove unwanted information, temporary files, your browsing history, huge log files and even the settings left behind by previously installed software.

A brand new CCleaner 6 feature is the Performance Optimizer, which can give your PC such an improvement, it can even be measured in benchmarks.

Another new feature of CCleaner 6 is the improved software updater. This helps to ensure that you always have the very latest versions installed.

It also includes security tools such as a data eraser so you can make sure important files and information is completely removed from your PC.

1 year - 1 PC
CCleaner 6
PRO
14 Powerful PC Optimization Tools
Boosts PC speed by up to 34%

EXTRA 20% OFF!

USE **PROMAR20** AT CHECKOUT

ALSO CONSIDER **Acronis True Image 2025 for $30** **SAVE 50%**

software.maximumpc.com

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

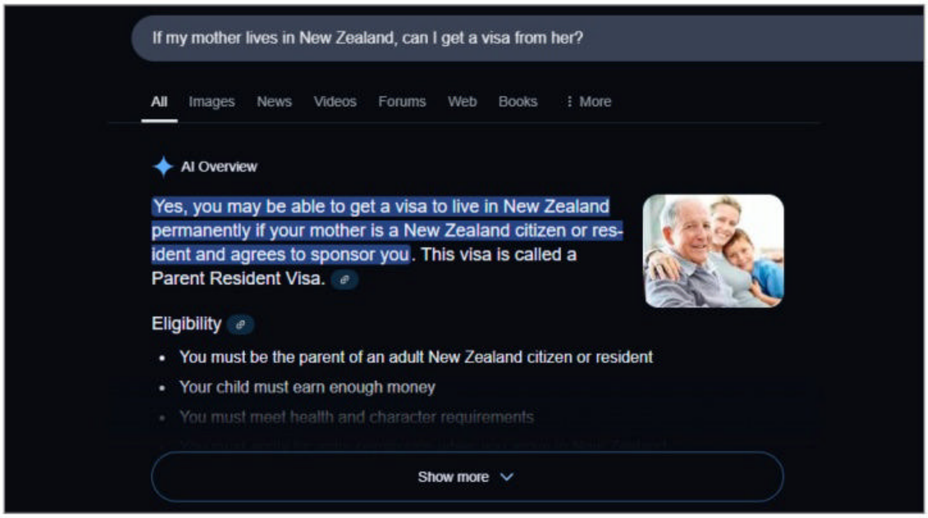

If my mother lives in New Zealand, can I get a visa from her?

Yes, you may be able to get a visa to live in New Zealand permanently if your mother is a New Zealand citizen or resident and agrees to sponsor you. This visa is called a Parent Resident Visa.

Eligibility

- You must be the parent of an adult New Zealand citizen or resident
- Your child must earn enough money
- You must meet health and character requirements

RESEARCH, RESEARCH, RESEARCH

Google has for the last few months integrated Gemini into its search queries. Got a question? Chuck it to Google, and Gemini will give you a response. Be wary, though; Gemini still prioritizes pretending to be a human over accurate information, plus some of the data it pulls can be erroneous, out-of-date, or just incorrect. It's worth searching a bit deeper and checking its answers, particularly if you're asking deeper questions.

MAKE – USE – CREATE

60
How to detect AI content

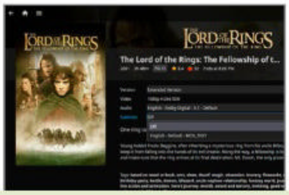

64
Build your own streaming server

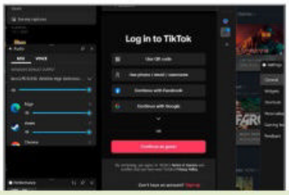

68
Beat Games with Edge Game Assist

© WIKIMEDIA, STABLE DIFFUSION, MP3TAG, JELLYFIN, MICROSOFT

ZAK STOREY
CONTRIBUTOR

COOLING ORIENTATION

I've been having a lot of fun lately on a custom build inside NZXT's H9 Elite chassis. It's a curious case; dual-chamber design, covered in glass—even the roof. NZXT's worked around that roof problem with some side ventilation, but it's still a bit of a hindrance for cooling solutions.

You've got three fans on the side, acting as an intake; three on the floor you can populate; three in the roof (opposed to that glass panel), and one at the rear.

I stacked it with an RTX 5090. plus an AMD Ryzen 9 9900X, combined with a 360mm AIO in the roof, along with six intakes and one rear exhaust. It got toasty during testing—my office hit 40 C.

I have found a solution, though. Usually, it's recommended that you install a 360mm radiator as an exhaust. With fishbowl designs, you typically have two sets of intakes: those three on the side, and the three in the floor.

I shifted the 360mm AIO to be side-mounted, getting access to room temperature air, letting the three 120mm fans in the roof exhaust without resistance from that thick radiator. The GPU is still actively cooled by its own dedicated 3x120mm intakes.

Problem solved; the GPU sits pretty at 78 C, and CPU temps have dropped by 10 C. And my office isn't giving me heatstroke.

↘ submit your How To project idea to: editor@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

Galaxy S25 Ultra

The pull-tab battery is the biggest advance in terms of repairability for the new Galaxy S25 Ultra.

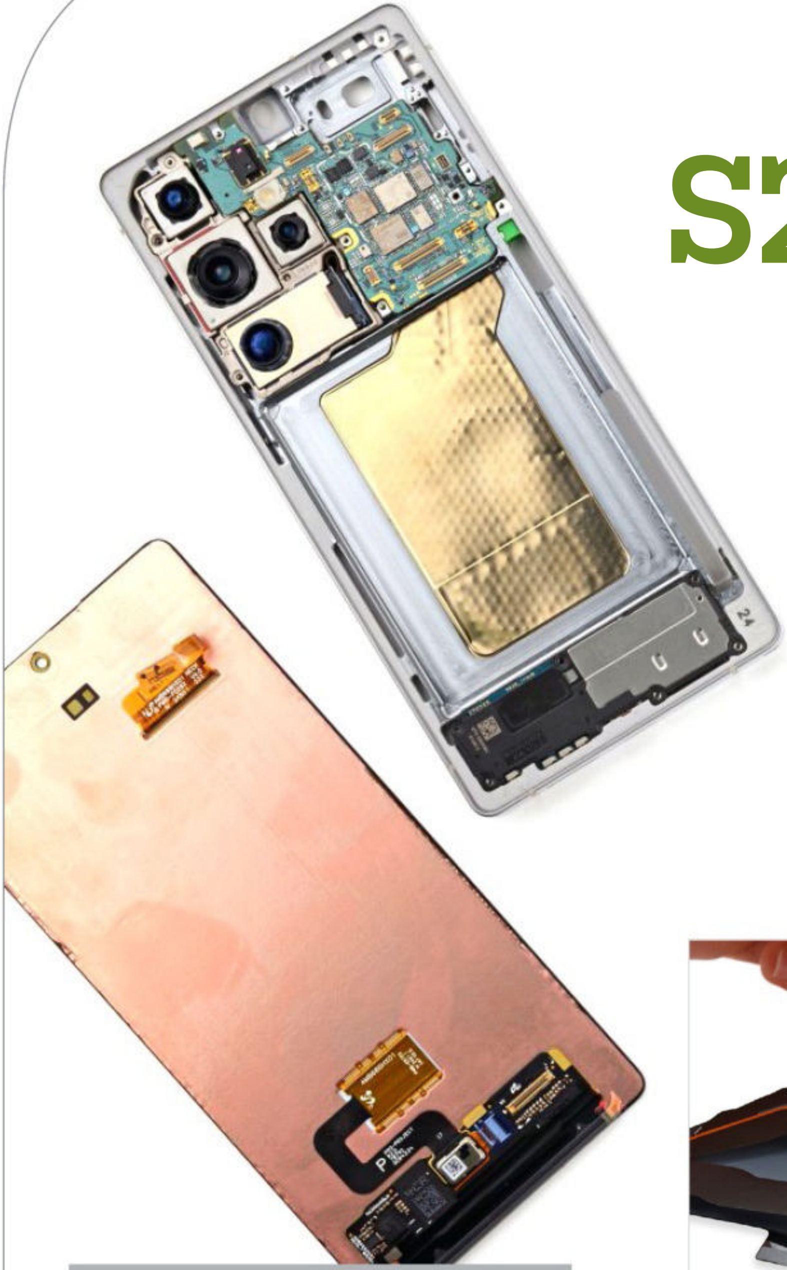

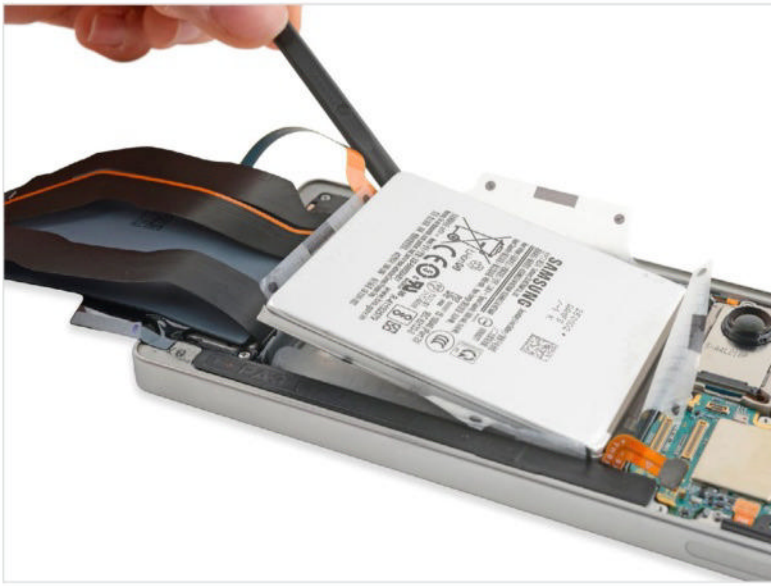

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products. To learn more, visit www.ifixit.com

At a glance, the S25 Ultra sheds some of its Galaxy identity with a decidedly Apple-inspired twist in the rounded corners of the frame.

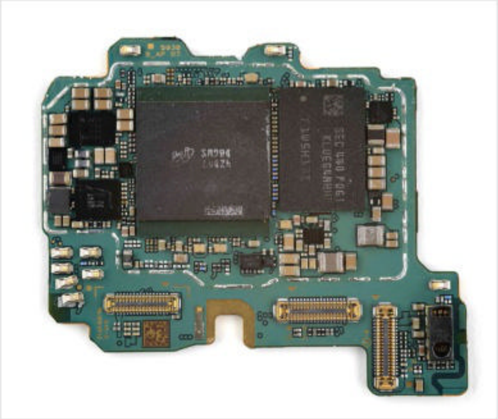

Inside, you'll find the cutting-edge Qualcomm Snapdragon 8 Elite chip with eight cores running at up to 4.47GHz.

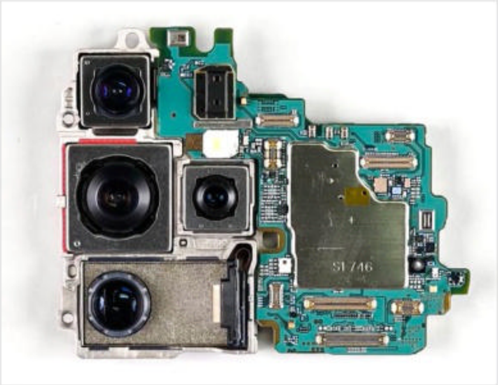

The modular rear camera array makes it possible to repair individual cameras rather than replacing the entire assembly.

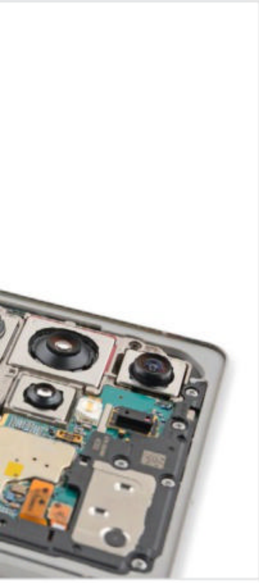

Overall, the Galaxy S25 Ultra is one of the most repair-friendly Galaxy phones in a decade.

Fans of Samsung's stylus are used to advanced features like Bluetooth connectivity. But in the Galaxy S25 Ultra, that wireless magic is gone.

BACKGROUND

When Samsung announced the Galaxy S25 Ultra, most of us expected the usual routine: better specs, minor design tweaks, and a new coat of paint to justify the next upgrade cycle. Instead, Samsung has delivered one of its most unexpected and intriguing flagships in years—at least where the phone's repairability is concerned.

MAJOR TECH SPECS

- 160 x 144 pixel 2.56-inch IPS backlit LCD
- 4.47GHz octa-core CPU
- 6.9-inch, 3,210 x 1,440, 120Hz AMOLED display
- Four rear cameras, 200MP, 50MP, 50MP, 10MP, 10x, 5x and 3x optical zoom, 8K60, 4K120 video recording
- 12MP front camera
- 12GB memory, 256GB storage
- 5G, WiFi 7, Bluetooth 5.4
- 5,000mAh battery
- 162.8 x 77.6 x 8.2mm, 218g
- From $1,299

KEY FINDINGS

- At a glance, the S25 Ultra sheds some of its Galaxy identity with a decidedly Apple-inspired twist in the rounded corners of the frame and the suspiciously iOS-like Quick Settings layout. As most tech reviewers have noted, though, the S25 Ultra is mainly ultra iterative over last year's flagship phone.

- However, building on the S23's pull tab mechanism that was designed to make battery replacement easier, this latest iteration of the pull tab is easier to operate still. As a result, the S25 Ultra could be one of Samsung's most repair-friendly phones in a long time.

- This time, Samsung has installed four neat sticky tabs that 'hug' the battery cell. Those tabs peel up easily, releasing the battery in a few seconds without an ounce of extra glue. You can pop it right out with no fuss, no mess, and no special solvents. For consumers and repair pros alike, this is a breath of fresh air.

- The S25 Ultra packs a multi-lens setup, including an upgraded 50 MP ultrawide sensor that promises improved image quality. Each camera module is independent, and can be removed or replaced separately. This means that if one lens suffers a nasty bump or the sensor malfunctions, you don't have to scrap the entire camera system. The selfie unit wasn't so lucky, though—it's strongly epoxied in place. While we were able to get ours out unscathed, our teeth were clenched the whole time.

- If some of the news is good, the broader picture isn't all sunshine. Official repair guides are lackluster, unintuitive, and complicated. Parts availability can be spotty and excruciatingly expensive, which undermines the concept of easily repairable hardware if you can't source the components needed for a fix.

- Repairability Score: 8 out of 10 (10 is easiest to repair). In the end, Samsung's Galaxy S25 Ultra manages to deliver a surprising bright spot for repair enthusiasts: the pull-tab battery alone is a huge leap. Combine that with modular cameras and the continued absence of enforced parts pairing, and you have one of the most repair-friendly Galaxy phones in a decade. ⏻

How to detect AI content

YOU'LL NEED THIS

EXIFTOOL
(Optional)

IN MARCH 2024, author Joanna Maciejewska tweeted: "I want AI to do my laundry and dishes so that I can do art and writing, not for AI to do my art and writing so that I can do my laundry and dishes." Her comments reflect the growing trend of using AI tools like LLMs (Language Learning Models) to generate text, images, and other types of media that mimic human creativity.

Many LLMs are trained on existing content without the consent of the creator, so they aren't always compensated properly for derivative works. In academic circles, bad actors sometimes pass off AI-generated content as their own to earn recognition and qualifications. What's even worse is that AI doesn't get things right all the time. Recent analysis of OpenAI's latest ChatGPT-4o model gives it an accuracy rating of 88 percent. If you're well-versed in a particular subject, this can be an excellent way to detect if text content is AI-generated.

In this guide, we'll explore how to identify inconsistencies that plague AI-generated content. You'll learn more about the tone and style of AI text, as well as telltale signs of generated images like flawed hands and backgrounds. We'll also touch on detecting AI-generated audio and videos. **–NATE DRAKE**

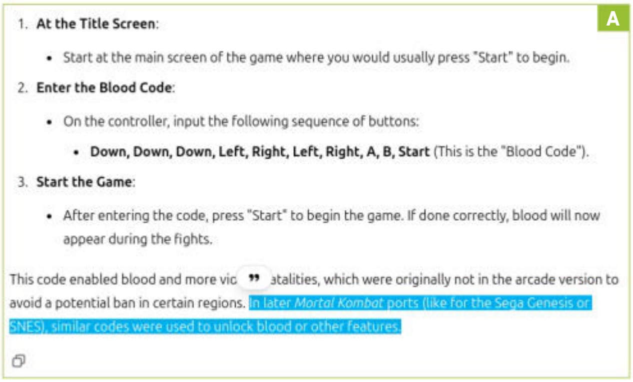

1. **At the Title Screen**:
   - Start at the main screen of the game where you would usually press "Start" to begin.
2. **Enter the Blood Code**:
   - On the controller, input the following sequence of buttons:
     - **Down, Down, Down, Left, Right, Left, Right, A, B, Start** (This is the "Blood Code").
3. **Start the Game**:
   - After entering the code, press "Start" to begin the game. If done correctly, blood will now appear during the fights.

This code enabled blood and more vio ❞ atalities, which were originally not in the arcade version to avoid a potential ban in certain regions. In later *Mortal Kombat* ports (like for the Sega Genesis or SNES), similar codes were used to unlock blood or other features.

1 RECOGNIZE FACTUAL ERRORS IN AI TEXT

One of the easiest ways to detect AI-generated text is to read through it carefully. LLMs work by predicting the next word in a sentence. This means the style tends to be more uniform and less prone to rare words than text written by a human.

» The same is true for careful writers—especially those who don't speak English as a first language. Bad actors may doctor AI-generated text to introduce more obscure terms or spelling/grammatical errors to make it harder to recognize. This may explain why a 2021 study using the GPT-3 model found that humans were unable to reliably detect AI text just by reading it.

» Still, LLMs can sometimes still give themselves away through inconsistencies in the text. These may be factual, or just abrupt shifts in the tone, style, or topic.

» For instance, in the example above, ChatGPT's GPT-4o model was asked to write instructions to activate the 'blood code' cheat in *Mortal Kombat* (1993). Experienced gamers will know that both Sega and Nintendo coded the home conversions of the game to be more family friendly by rendering blood in gray.

» However, the Genesis version of the game contained a secret code to restore the blood to its natural color. During our tests, ChatGPT incorrectly stated that this cheat also existed in the original arcade and SNES versions of the game. The button sequence listed is also designed to invoke a cheat menu—not the 'blood code' itself. [**Image A**]

» Detecting AI usage in this way requires careful fact-checking and research of text content. Human beings are also prone to factual errors, so this alone shouldn't be taken as evidence that an LLM has generated the text.

2 LINGUISTIC ANOMALIES IN LLMS

Language models are trained on existing human-written text, so sometimes it's hard to tell the difference between AI and an uninspired writer. In the context of tech, for instance, we've lost count of the number of articles that begin, 'In today's interconnected world...'

» This kind of blandless can also sometimes spill into the grammar of AI text. The below paragraph was generated by DeepSeek's DeepThink R1 language model in response to a prompt for an article on the advantages of switching to Linux:

» 'One of the biggest advantages of Linux is its cost. Unlike other operating systems, Linux is completely free to download and use. You can say goodbye to expensive licenses and subscriptions, and still enjoy a powerful system that meets your needs.'

» In this case, the length of the sentences are short and uniform. In fact, they're each 10-15 words long. They also all follow the SVO (Subject-Verb-Object) word order.

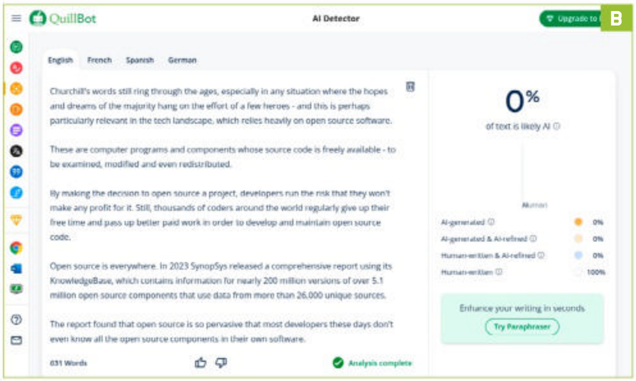

Churchill's words still ring through the ages, especially in any situation where the hopes and dreams of the majority hang on the effort of a few heroes - and this is perhaps particularly relevant in the tech landscape, which relies heavily on open source software.

These are computer programs and components whose source code is freely available - to be examined, modified and even redistributed.

By making the decision to open source a project, developers run the risk that they won't make any profit for it. Still, thousands of coders around the world regularly give up their free time and pass up better paid work in order to develop and maintain open source code.

Open source is everywhere. In 2023 SynopSys released a comprehensive report using its KnowledgeBase, which contains information for nearly 200 million versions of over 5.1 million open source components that use data from more than 26,000 unique sources.

The report found that open source is so pervasive that most developers these days don't even know all the open source components in their own software.

631 Words

0% of text is likely AI

AI-generated 0%
AI-generated & AI-refined 0%
Human-written & AI-refined 0%
Human-written 100%

© OPENAI, QUILLBOT

A human being is more likely to mix these things up by varying word order and sentence length throughout the text.

» Another key indicator of AI text can be when the tone is rather sterile, without any strong opinions of personal anecdotes. This said, LLMS can be requested to use a more conversational tone in writing. When we did this for the above paragraph with DeepThink R1 it read:

» 'First off, let's talk about cost. Linux is free. Yep, you read that right. No expensive licenses, no subscriptions, no sneaky fees. You can download it, install it, and start using it without spending a dime.'

» As readers will see, the text is still written in very brief staccato sentences. The LLM also makes use of repetition, eg. the word 'no'. Although the tone is considerably more casual, it also still lacks personal insights and unique perspectives that a human can provide.

» AI also sometimes gives itself away through use of region-specific language. In this example, the text uses the phrase 'without spending a dime'. This is a regular expression in American English. However, it would be unusual for, say, an Australian to use it when writing for people in their home country.

3 USE AI DETECTION TOOLS

If you're considering using automated tools to detect AI, we recommend generating text with a similar theme to what you're currently investigating. Ideally, you should also test with some text that you know has been written solely by a human.

» Most tools like GPTZero (https://app.gptzero.me) allow you to paste text into a box. They will then display the likelihood (displayed as a percentage) that the text was written by AI, human beings, or a mixture of the two. [**Image B**]

» For instance, after running our 'conversational' Linux piece generated by DeepThink R1 through GPTZero's 'basic' scan, it told us there was a 100 percent probability that the text was AI generated. Signing in with our Google Account allowed us to run an 'Advanced Scan', which expanded further on 'high impact' sentences that had a high probability of being AI.

» GPTZero and other AI detection tools warn that these methods are probabilistic and predictive. In other words, there's no way to be certain. That said, if enough AI-like vocabulary or wording appears in text, it's likely that an LLM had a hand in creating it.

DETECTING VIDEO DEEPFAKES

In 2017, the 'Synthesizing Obama' program modified footage of the former President's speeches to a different audio track. The techniques used to adapt video of a person's face is considered to be one of the earliest 'deepfakes', though the techniques used weren't all AI-based.

Though this was done for academic research, deepfake videos are often used for fake news or for bogus celebrity sex tapes. High-quality deepfakes deploy AI to train on hours of genuine videos of the target. This has been made easier with the rise of open-source tools like Faceswap.

Most of these use an 'autoencoder'. This reduces images to a lower dimensional latent space, then derives information like facial features, body type, and posture. These can be coupled with generative adversarial networks (GANs), which use algorithms to see if a frame is generated, then modify it. This can make deepfake videos look extremely realistic.

Fortunately, most deepfake efforts are created by amateurs. As with AI-generated images, the lighting can seem off. One of the most common deepfake techniques is to superimpose the face of a celebrity onto an actor's body. Look for blurring and color/texture differences between the face and neck. If the person is speaking, listen to ensure that their lip movements sync exactly with the audio.

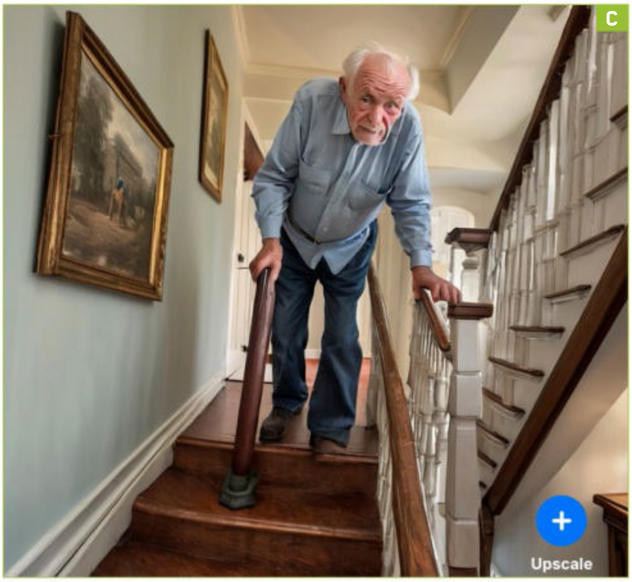

© WIKIMEDIA, STABLE DIFFUSION

» If you're still uncertain, run the text through multiple AI detection tools. For example, QuillBot (https://quillbot.com/ai-content-detector) also detected that our DeepSeek-generated text was 100 percent AI content. It also correctly detected that a TechRadar article written by Nate on a similar theme contained 0 percent AI content.

» Unfortunately, a quick search of sites like Reddit shows that many innocent students have been unfairly flagged by AI detection tools. One of the best ways to find out for sure is to inspect the document history in the word processor originally used to create the text.

» For example, in Google Docs users can hold Ctrl-Alt-Shift-H to view the version history of a document. A piece of text written by a human will list multiple versions and edits. If text is AI-generated, it's likely to be pasted in one block.

4 ANALYZING AI-GENERATED IMAGES

In the words of Sherlock Holmes, circumstantial evidence is occasionally very convincing. In the context of AI, this means generated images can be much

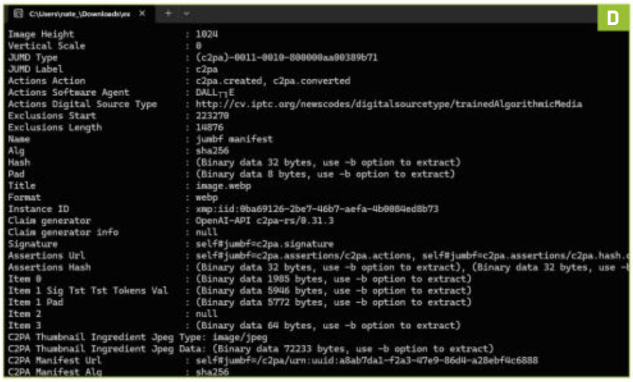

```
Image Height                    : 1024
Vertical Scale                  : 0
JUMD Type                       : (c2pa)-0011-0010-800000aa00389b71
JUMD Label                      : c2pa
Actions Action                  : c2pa.created, c2pa.converted
Actions Software Agent          : DALL┬┐E
Actions Digital Source Type     : http://cv.iptc.org/newscodes/digitalsourcetype/trainedAlgorithmicMedia
Exclusions Start                : 223270
Exclusions Length               : 14876
Name                            : jumbf manifest
Alg                             : sha256
Hash                            : (Binary data 32 bytes, use -b option to extract)
Pad                             : (Binary data 8 bytes, use -b option to extract)
Title                           : image.webp
Format                          : webp
Instance ID                     : xmp:iid:0ba69126-2be7-46b7-aefa-4b0884ed8b73
Claim generator                 : OpenAI-API c2pa-rs/0.31.3
Claim generator info            : null
Signature                       : self#jumbf=c2pa.signature
Assertions Url                  : self#jumbf=c2pa.assertions/c2pa.actions, self#jumbf=c2pa.assertions/c2pa.hash.
Assertions Hash                 : (Binary data 32 bytes, use -b option to extract), (Binary data 32 bytes, use -
Item 0                          : (Binary data 1985 bytes, use -b option to extract)
Item 1 Sig Tst Tst Tokens Val   : (Binary data 5946 bytes, use -b option to extract)
Item 1 Pad                      : (Binary data 5772 bytes, use -b option to extract)
Item 2                          : null
Item 3                          : (Binary data 64 bytes, use -b option to extract)
C2PA Thumbnail Ingredient Jpeg Type: image/jpeg
C2PA Thumbnail Ingredient Jpeg Data: (Binary data 72233 bytes, use -b option to extract)
C2PA Manifest Url               : self#jumbf=/c2pa/urn:uuid:a8ab7da1-f2a3-47e9-86d4-a28ebf4c6888
C2PA Manifest Alg               : sha256
```

easier to detect than text. For example, a web user who stumbles across a color photo of Abraham Lincoln playing an electric guitar doesn't need any further evidence to know it's fake.

» Unfortunately, not all AI images are quite so easy to spot. If an image contains people, first inspect the hands. AI is still notoriously bad at rendering these correctly, so check carefully for extra/missing fingers, superfluous joints, or veiny palms. Many generators also struggle with hair. Look for hair that sits too far above the head, or even blends into clothing.

» Unless specifically prompted otherwise, many AI image generators will also blur the background to prevent obvious visible inconsistencies. If the background is clear, check any architecture carefully. Walls, ceilings, and stairs can be curved or misaligned. Rendered objects can sometimes appear much larger or smaller than they would in real life.

» AI also often struggles with rendering accessories correctly. For instance, supposedly matching earrings may appear in different sizes. Some might look to be fused into clothing or in the wrong place. All of these incongruities are present in the image we generated in Stable Diffusion of an old man walking up stairs. There are two sets of banisters, and his cane is overly large. He also wears a belt over an untucked shirt. [**Image C**]

» Even if there are no obvious mistakes, objects and people in the foreground of AI images tend to have a sheen to them. The appearance is overly rendered, similar to how characters appear in modern video games.

» Text in AI images can also be a giveaway. In real life, text is usually written clearly in a consistent size. Check for spelling errors or misshapen letters.

5

ADVANCED AI IMAGE DETECTION

If a picture contains no obvious inconsistencies, check if it's legitimate by using Google or TinEye to run a reverse image search. If it has been published by reputable outlets, then it's more likely to be real.

» You can use utilities like ExifTool (https://exiftool.org) to view more detailed embedded information. For an image that was taken by a human, this is likely to list realistic information such as the camera model or where the photo was taken.

» If you're running Windows, once exiftool(-k).exe is extracted, you can list metadata by dragging and dropping an image file onto the executable. If you do this for AI images, this can reveal clues to their origins. For instance, an image we'd generated contained Coalition for Content Provenance and Authenticity (C2PA) data when analyzed in ExifTool. This is designed to record information about the provenance of an image. It listed DALL-E as its 'Actions Software Agent'. It also included a link to OpenAI's API.

» Remember, image metadata is relatively easy to edit or strip from an image, so you shouldn't rely on utilities like ExifTool alone to determine if an image is genuine. [**Image D**]

» Some image generators also output pictures in predictable ways. For instance, by default, DALL-E 3 only supports resolutions of 1024x1024, 1024x1792, or 1792x1024. Developers who generate images using DALL-E's API can use other resolutions. Stable Diffusion defaults to 512x512 images, though it can support resolutions of up to 2048x2048.

» There are a number of guides on how to differentiate between real and AI images based on their entropy. This works by calculating the local entropy or randomness for each RGB (Red, Green, and Blue) channel in an image, then the 'Shannon entropy' for surrounding pixels.

» Traditionally, AI images had lower entropy due to having less complex textures and shapes. However, Modern AI generators can create high-entropy images, so this isn't a reliable indicator.

6

AI IMAGE DETECTION TOOLS

As with text detection tools, there are sites that claim to be able to spot the difference between real and AI-generated images. Most of these work by analyzing image pixels to detect unusual configurations. They can also examine existing data sets to see if any people or objects in the image match existing AI models.

» Popular options include 'AI or Not' (www.aiornot.com) and Illuminarty (https://app.illuminarty.ai). Whichever tool you choose, the site will prompt you to upload the image. You'll be shown the likelihood (as a percentage) that the image was AI or human-generated.

» Even the best tools can't always reliably detect AI images. The two we've mentioned have an accuracy rate of around 85 and 71 percent respectively. Some tools will also show only limited information unless you sign up and/or have a paid subscription.

» If you're running a Chromium-based browser like Chrome, you can also install the 'Hive AI Detector' [**Image E**] extension, which can check if any in-browser content is AI generated. Once installed, right-click the image, video, or audio in question, and select 'Hive AI Detector' to view its AI content analysis.

7

ANALYZING AI AUDIO

In September 2024, a YouTube video appeared entitled 'Best of 1950s to 1970s Christmas Carols'. Despite ostensibly containing festive classics sung by the likes of Frank Sinatra, the compilation video was quickly outed as AI 'slop'.

» The first indication of this was around 36 minutes in, where 'O Little Town of Bethlehem' was sung to the tune of 'Silent Night'. This didn't stop the video from amassing over 5 million views and almost 3,000 comments.

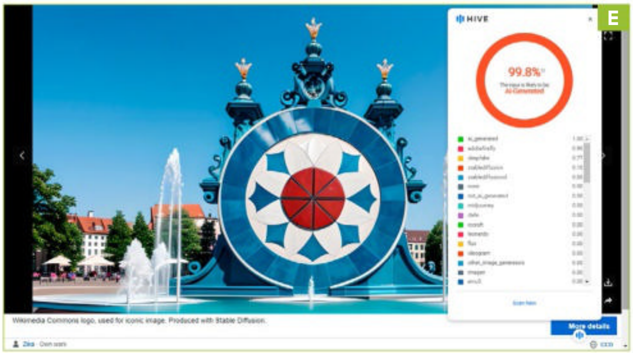

Wikimedia Commons logo, used for iconic image. Produced with Stable Diffusion.

© WIKIMEDIA/HIVE AI, PHIL HARVEY(EXIFTOOL DEVELOPER)

ANATOMY OF AN AI IMAGE

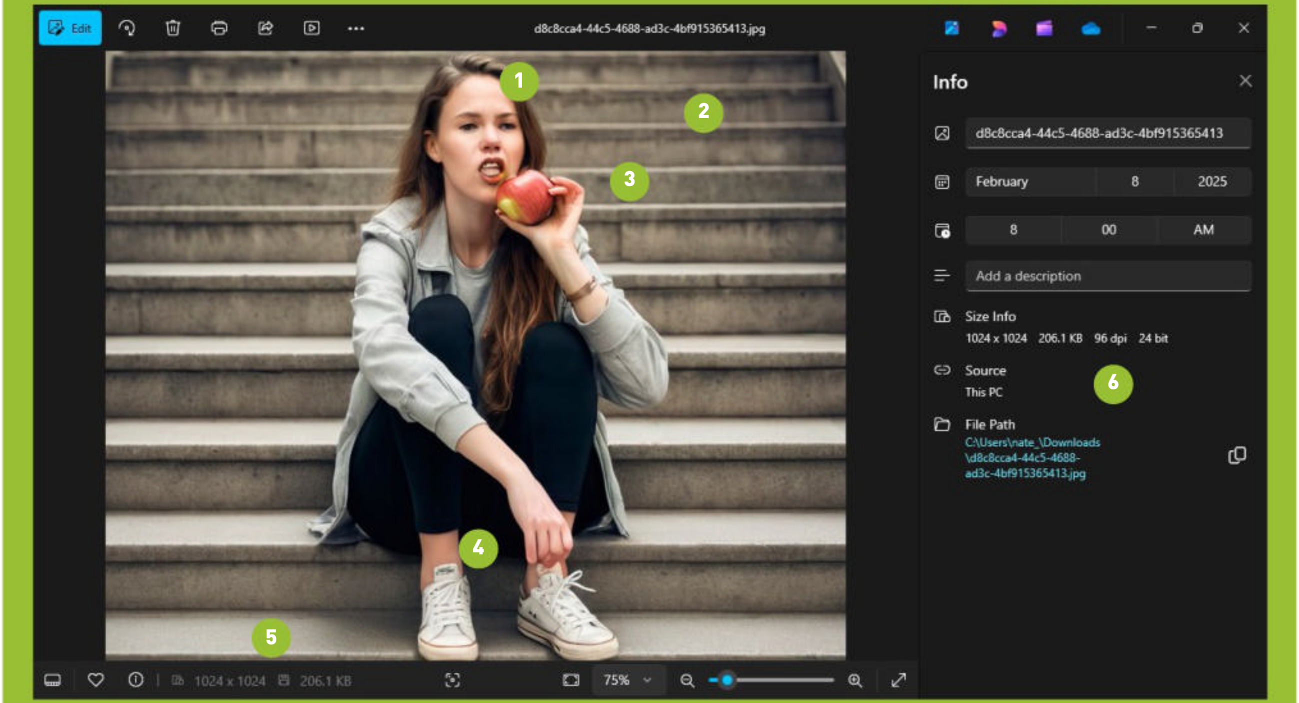

Info

d8c8cca4-44c5-4688-ad3c-4bf915365413

February 8 2025

8 00 AM

Add a description

Size Info
1024 x 1024 206.1 KB 96 dpi 24 bit

Source
This PC

File Path
C:\Users\nate_\Downloads\d8c8cca4-44c5-4688-ad3c-4bf915365413.jpg

1. AI SHEEN
Objects and people in the foreground of AI images often have an unusual sheen caused by inconsistent lighting. They can also appear much more detailed than background objects.

2. BLURRED BACKGROUNDS
Without specific prompts, most AI Image generators will display a blurred and/or very uniform background. Objects in the background can also appear indistinct. The architecture may sometimes appear warped or inconsistent.

3. HUMAN HANDS
AI Image Generators still often struggle to render hands correctly. Fingers may also appear to be fused together (as is the case here). Additionally, there may be extra/missing fingers or joints.

4. BLURRED TEXT
Although some generators are getting better at this, text in AI images can appear distorted. Letters may be warped or of different sizes. In this case the other sneaker also has no lettering.

5. SET RESOLUTIONS
Certain AI image generators default to fixed resolutions. For instance, DALL-E 3 creates 1024 x 1024 images. Skilled programmers can utilize the generator's API to create pictures with different sizes.

6. IMAGE METADATA
You can use utilities like ExifTool to examine an image's metadata. Many AI image generators record their origins here. Real images can also include EXIF data like the location camera model.

» Shortly after, a thread sprung up on Reddit of supposed covers of Christmas songs by the likes of 'Dean Snowfield' cropping up on sites like YouTube and Spotify.

» These fake AI artists were quickly detected, as their bios don't tend to be fleshed out, or contain only short AI-generated phrases. Although supposedly by different singers, some of the voices were eerily similar to each other or famous artists like Nat King Cole.

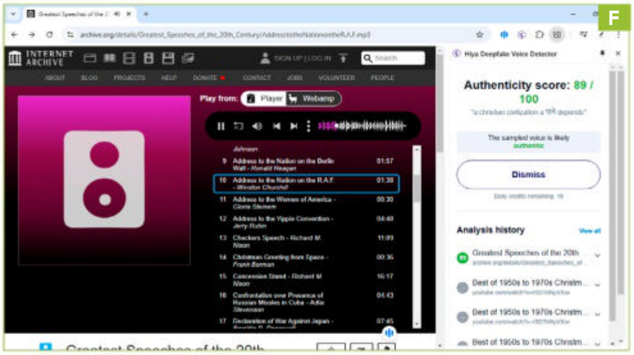

Authenticity score: 89 / 100

The sampled voice is likely authentic

Analysis history

Greatest Speeches of the 20th

Best of 1950s to 1970s Christm...

Best of 1950s to 1970s Christm...

Best of 1950s to 1970s Christm...

9 Address to the Nation on the Berlin Wall - Ronald Reagan 01:57
10 Address to the Nation on the R.A.F. - Winston Churchill 01:38
11 Address to the Women of America - Gloria Steinem 00:30
12 Address to the Yippie Convention - Jerry Rubin 04:48
13 Checkers Speech - Richard M. Nixon 11:09
14 Christmas Greeting from Space - Frank Borman 00:36
15 Concession Stand - Richard M. Nixon 16:17
16 Confrontation over Presence of Russian Missiles in Cuba - Adlai Stevenson 04:43
17 Declaration of War Against Japan - 07:45

© CG SECURITY, INTERNET ARCHIVE/HIYA

» During our tests, Hive AI Detector had no joy with streaming videos on YouTube. However, 'Hiya Deepfake Voice Detector' [Image F] is also available in the Chrome Web Store. You'll need to register with your email.

» Once the extension is installed, return to the tab where the audio is playing, then click to open the extension. Choose 'Start analyzing'. While this worked well for videos containing speech, the extension couldn't analyze the voices singing in the Christmas carol video.

» Although AI audio can be trained to mimic existing artists, it can struggle with words or phrases that don't appear in its dataset. This means you may find that AI fakery will slur or mispronounce certain words.

» Pay attention to intonation. If a voice sounds toneless or unemotional, notably when describing something that normally provokes feelings, it may be AI-generated.

Build your own streaming server

YOU'LL NEED THIS

JELLYFIN
(https://jellyfin.org) comes as a server component for PC and Mac, plus client devices across all major mobile, desktop, and streaming box/TV platforms

IF YOU PREFER OWNED MEDIA to paying expensive monthly subscription services for the likes of Netflix or Amazon, then why not have the best of both worlds? Self-hosted services like Plex and Jellyfin make it possible to present your media in a way that's easy and quick to access—no navigating multiple platforms or disc menus. Plex may be the king, but many of its best features are hidden away behind the Plex Pass paywall.

Thankfully, you can enjoy virtually all of Plex's best features without paying a dime, thanks to the brilliant open-source Jellyfin. Since forking from another media server (Emby), Jellyfin has gone from strength to strength, and is now our go-to self-hosted media streamer. Here, we'll show you how to get it installed, add your libraries, and get to grips with its navigation and playback controls. We'll also introduce you to its powerful dashboard so you can explore what it offers even further. **–NICK PEERS**

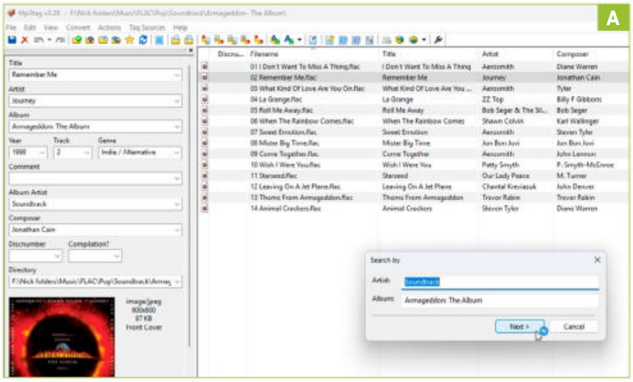

Title: Remember Me
Artist: Journey
Album: Armageddon: The Album
Year: 1998
Track: 2
Genre: Indie / Alternative
Album Artist: Soundtrack
Composer: Jonathan Cain

Filename	Title	Artist	Composer
01 I Don't Want To Miss A Thing.flac	I Don't Want To Miss A Thing	Aerosmith	Diane Warren
02 Remember Me.flac	Remember Me	Journey	Jonathan Cain
03 What Kind Of Love Are You On.flac	What Kind Of Love Are You ...	Aerosmith	Tyler
04 La Grange.flac	La Grange	ZZ Top	Billy F Gibbons
05 Roll Me Away.flac	Roll Me Away	Bob Seger & The Sil...	Bob Seger
06 When The Rainbow Comes.flac	When The Rainbow Comes	Shawn Colvin	Karl Wallinger
07 Sweet Emotion.flac	Sweet Emotion	Aerosmith	Steven Tyler
08 Mister Big Time.flac	Mister Big Time	Jon Bon Jovi	Jon Bon Jovi
09 Come Together.flac	Come Together	Aerosmith	John Lennon
10 Wish I Were You.flac	Wish I Were You	Patty Smyth	P. Smyth-McEnroe
11 Starseed.flac	Starseed	Our Lady Peace	M. Turner
12 Leaving On A Jet Plane.flac	Leaving On A Jet Plane	Chantal Kreviazuk	John Denver
13 Theme From Armageddon.flac	Theme From Armageddon	Trevor Rabin	Trevor Rabin
14 Animal Crackers.flac	Animal Crackers	Steven Tyler	Diane Warren

Search by
Artist: Soundtrack
Album: Armageddon: The Album

1 PREPARE YOUR MEDIA

Before diving into installing Jellyfin, make sure your owned media is suitably named and organized so it can be easily imported into one of Jellyfin's libraries. You'll want to split different types of media into different categories, from simple libraries spanning all your movies, music, and TV, to potentially splitting out larger collections across multiple libraries (TV—Drama, TV—Comedy, and so on).

» There's an art to naming your files, so Jellyfin can identify them and do the hard work of scraping various online metadata databases to provide you with titles, synopses, artwork, and more. Jellyfin follows Plex and other media servers, so it's reasonably intuitive—visit the Media section at https://jellyfin.org/docs for details. However, the following conventions should cover most bases:

» MOVIES—name your movie Film Title (1990).mp4, and place it inside a folder of the same name, So Film Title (1990)\Film Title (1990).mp4. Extras should be organized into their own subfolders like featurettes, deleted scenes, trailers, and so on inside the main Film Title (1990) folder.

» TV SHOWS—organize these into Show Title\Season 1\Show Name—S01E01.mp4, and so on. You can append the episode title on the end for easier visual identification of each file. Use the excellent Rename My TV Series app (www.tweaking4all.com/home-theatre/rename-my-tv-series-v2) to speed up the process of making your TV shows ready for Jellyfin.

» MUSIC—Follow the classic structure utilized by the likes of Windows Media Player Legacy and other CD ripping tools: Artist\Album Title\01.Track name.flac, and so on. Tools like MP3tag (www.mp3tag.de/en) as shown in [**Image A**] and MusicBee (https://getmusicbee.com) can help tidy up messy libraries by embedding tags and artwork into the files themselves. Check out MetaX ($9.95, www.danhinsley.com) if you want to embed tags into your video files, too.

2 INSTALL JELLYFIN

The server portion of Jellyfin is available to install on Linux, Windows, and macOS as native apps—head over to https://jellyfin.org/downloads for direct downloads. Windows users should choose the AMD64 build for x86/x64 chips.

» You can also install it as a Docker or Podman container—a good choice for a low-powered server that's running 24-7. Visit https://jellyfin.org/docs/general/installation/container#docker for a detailed set of instructions covering both platforms.

» Installation in Windows is simple: just double-click the .exe file and follow the prompts—if you get a SmartScreen warning, click More info, followed by Run anyway. The recommended Basic Install is the simplest procedure to follow, but note that Jellyfin will only run when your PC is logged in to your user account. Although Service Install is possible—basically running Jellyfin in the background regardless of who's logged in—it's not

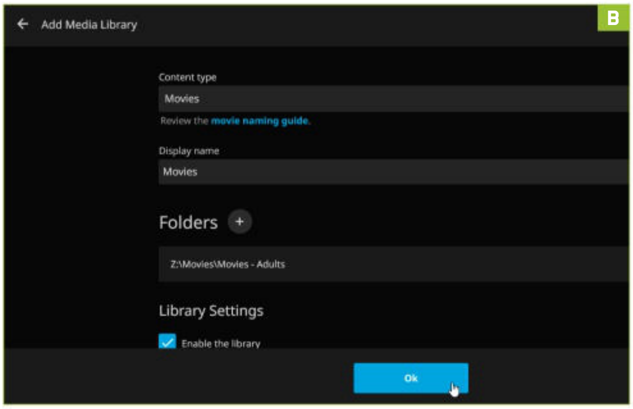

Content type: Movies
Review the movie naming guide.
Display name: Movies
Folders: Z:\Movies\Movies - Adults
Library Settings
☑ Enable the library

© MP3TAG, JELLYFIN

ACCESS YOUR MEDIA

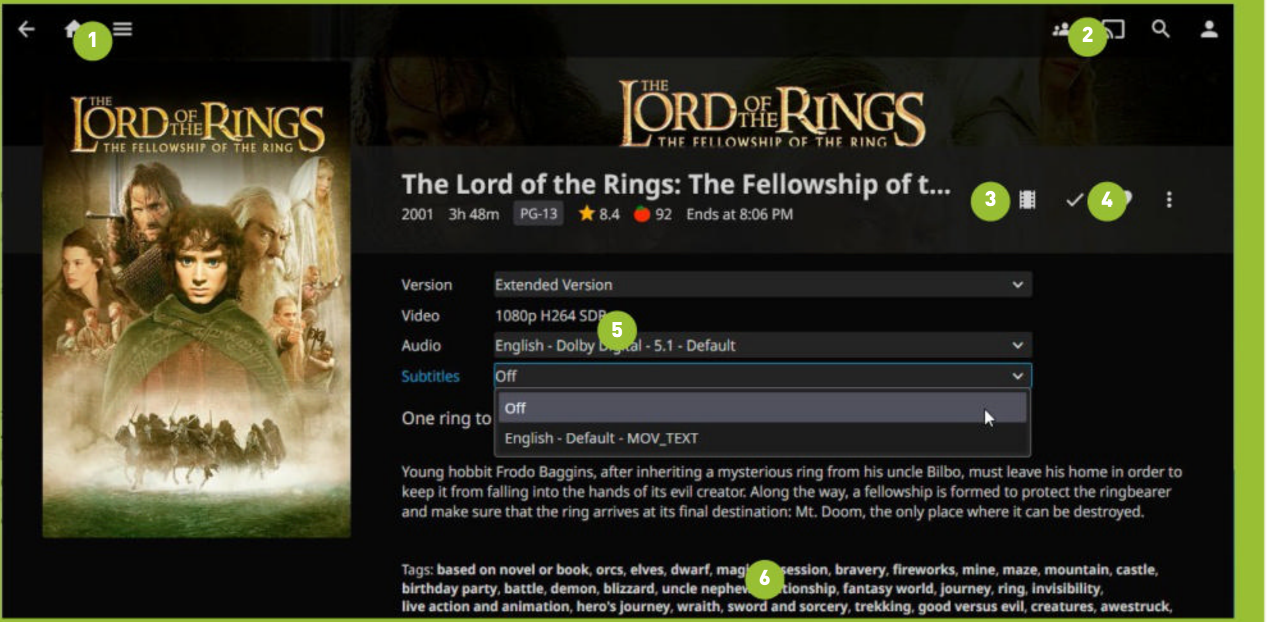

The Lord of the Rings: The Fellowship of t...

2001 3h 48m PG-13 ★ 8.4 92 Ends at 8:06 PM

Version: Extended Version
Video: 1080p H264 SDR
Audio: English - Dolby Digital - 5.1 - Default
Subtitles: Off

- Off
- English - Default - MOV_TEXT

One ring to

Young hobbit Frodo Baggins, after inheriting a mysterious ring from his uncle Bilbo, must leave his home in order to keep it from falling into the hands of its evil creator. Along the way, a fellowship is formed to protect the ringbearer and make sure that the ring arrives at its final destination: Mt. Doom, the only place where it can be destroyed.

Tags: **based on novel or book, orcs, elves, dwarf, magi ession, bravery, fireworks, mine, maze, mountain, castle, birthday party, battle, demon, blizzard, uncle nephe tionship, fantasy world, journey, ring, invisibility, live action and animation, hero's journey, wraith, sword and sorcery, trekking, good versus evil, creatures, awestruck,**

1. NAVIGATION
Click the hamburger icon to reveal a slide-out menu with shortcuts to other media libraries, your media administration, and your user profile.

2. ADVANCED PLAYBACK
Jellyfin supports SyncPlay, a way for multiple users with access to your server to watch content in real-time together—from different devices, rooms, or even locations.

3. PLAY MOVIE
Click the play button to start playing your movie. If you've downloaded or added trailers to your library, you can play them by clicking the cinematic film icon.

4. MARK FILMS
Use the tick button to mark the film watched—it'll turn red. Click it again to revert to unwatched. The heart button lets you mark it as a favorite, accessible via Favorites on your home screen.

5. PLAYBACK OPTIONS
If you have multiple versions of a movie, you can switch between them using the Versions drop-down. Once selected, you can then switch between audio track and enable or disable subtitles as required.

6. TAGS
These tags are added automatically when the film is added to your library. Edit them in the Metadata Manager, or click one to view a list of all titles containing that tag.

recommended. First, you won't be able to add folders stored on network drives to your library, and it has also been known to impact attempts to use hardware acceleration.

» With that in mind, leave Basic Install selected, and click Next. Leave both default locations as they are (C:\Program Files and C:\ProgramData respectively) and click Next again. Finally, review a summary of your choices, and click Install to complete the installation process.

3 FIRST-RUN WIZARD

If you follow the basic install method, you'll need to open the newly added Jellyfin Tray App. This will automatically start Jellyfin in the background, and set it to automatically start each time you log into Windows. The tray icon can be used to change this behavior, plus start and stop Jellyfin—simply right-click the Jellyfin icon in your Taskbar's notification area. You can also select 'Open Jellyfin' to open the server's web-based interface in your browser, or go directly to http://localhost:8096 to start configuring your new media server.

» The wizard is self-explanatory—start by selecting your language (English is of course the default). Next, create a username and password for yourself as the server's master administrator—if you plan to make your server accessible over the internet, take the time to choose a strong password with the help of your password manager. Click Next again.

» Your next task is to set up your first media library. Click the Add Media Library button to get started. You'll be prompted to give it a suitably descriptive name and choose its content type: Movies, Music, Shows (TV), Books, Home Videos, and Photos, Music Videos, and Mixed Movies, and Shows are all available, but in most cases, you'll be focused on the first three categories.

» Although libraries can encompass multiple folders—just keep clicking + next to Folders to browse and select each one in turn—a smart move is to set up individual parent folders for each library. If your folders are stored on a network drive, the simplest way to get Jellyfin to recognize them is by assigning the network drive its own letter—browse to the drive (or shared folder) in File Explorer, right-click it, and choose 'Map network drive' to do so.

4 LIBRARY SETUP

After selecting a library type as shown in [Image B], you'll unlock a host of other options to work

WATCH ON ANY DEVICE

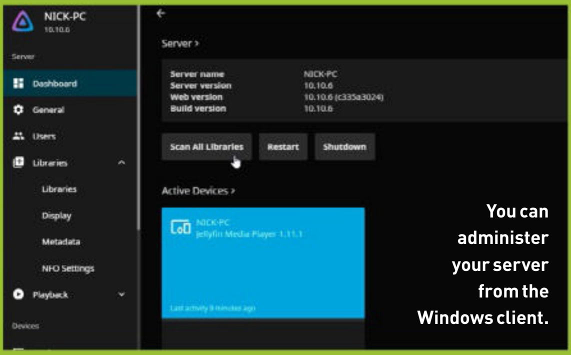

Server name NICK-PC
Server version 10.10.6
Web version 10.10.6 (c335a3024)
Build version 10.10.6

You can administer your server from the Windows client.

While you can access content through your web browser, you'll probably want to widen your options to include mobile devices, other computers, and even smart TVs. In recent years, Jellyfin has widened its reach to offer a wide range of clients—all of which are free to use, unlike its rivals.

You'll find native Jellyfin clients for all major platforms—search the relevant app store on your phone, tablet, streaming box (including Apple TV, Android TV, and Roku), or TV to find both the official app and several unofficial ones.

Alternatively, visit https://jellyfin.org/downloads/clients for a list, including the official desktop client for Windows, Mac, and Linux. This media player can also be used to administer your server—it's basically a wrapper for the web interface. One platform missing from the list is the Tizen-powered Samsung TV—visit https://github.com/jellyfin/jellyfin-tizen for details of sideloading the app on to your TV.

You can also access your Jellyfin music collection through dedicated music clients—our current favorite is Finamp (https://github.com/jmshrv/finamp), which in addition to offering a beautiful frontend for navigating and playing your library allows you to download portions of your library for offline listening.

through. First, make sure 'Enable the library' remains checked. You can leave 'Preferred download language' blank (Jellyfin will take its cues from your default language settings), but take the time to select United States under Country/Region. If your content has been tagged with a title tag, check 'Prefer embedded titles over filenames' for those titles that don't show up in any metadata provider databases.

» Work through the rest of the library options with the aid of the descriptions beneath each option. Options to look out for include to save downloaded artwork into the media folders (allowing you to easily find and edit them later if required), alongside an option to store individual items' metadata in separate .Nfo files alongside the media itself. A`lso, take a close look at the Trickplay and Chapter Image sections, which help generate useful thumbnails to help you navigate through videos by scrubbing or chapter selection. These can be resource and storage intensive, so use with care. Once complete, click OK—you can edit these options later via the Dashboard, so nothing is set in stone.

» After setting up your first library, you can create more—or add them later if you prefer. After clicking Next you'll be prompted to set your preferred metadata language and country/region, which will mimic whatever options you chose in the previous step. Then, you'll be given the opportunity to set up remote access, allowing you (and others) to connect to your Jellyfin server from outside your home. If you want to restrict access to only those devices on your local network, uncheck 'Allow remote connections to this server' before clicking Next followed by Finish.

5 SIGN IN AND BROWSE

You'll now be taken to the main Jellyfin sign-in page—enter the username and password you set up in step 3 and click 'Sign In'. Leave 'Remember me' checked if you're the only user of your PC to skip this step in future. You'll find yourself on the Jellyfin home page, which will display thumbnails of each library you've set up under 'My media'. Beneath this, you'll see newly added content listed in rows, one per library. Going forward, the home page will also offer convenient shortcuts to the media you're currently watching, as well as future episodes of shows you're watching under 'Continue watching' and 'Next up' respectively.

» Click on one of your library thumbnails to access that specific library, where you'll find a grid-like display of all the movies, shows, or artists in that collection. Above these thumbnails are a row of buttons to aid navigation—by default, 100 items are shown on each page, so use the arrow buttons to move between them (or note the A-Z list down the right-hand side, allowing you to browse your collection even more quickly). Next to this is a shuffle button—clicking this plays one title randomly from the library.

» The button to the right of shuffle lets you view your library in several different ways as seen in [Image C] : List displays movies in rows, including key information like video length, user rating, and whether a title has subtitles or not, while Banner and Thumbnail provide alternatives to the default Poster view—of the two, Thumbnail is more effective. Following on is the Sort button, letting you rearrange the order of your library's content any number of ways. Finally, things get interesting with the Filters button, which allows you to show only those titles matching a wide range of criteria, from movie rating and release year to specific tags and video type.

6 SELECT A TITLE

Sometimes, Jellyfin gets confused over identifying titles (particularly if you've misspelled it or given it the wrong year). As a result, you may find some are incorrectly listed in your library. The fix is

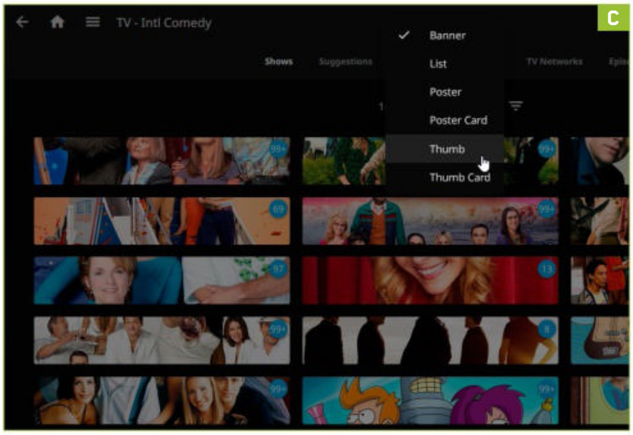

© JELLYFIN, MICROSOFT

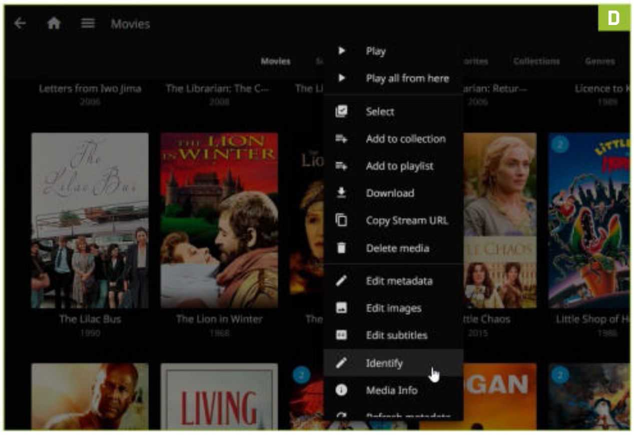

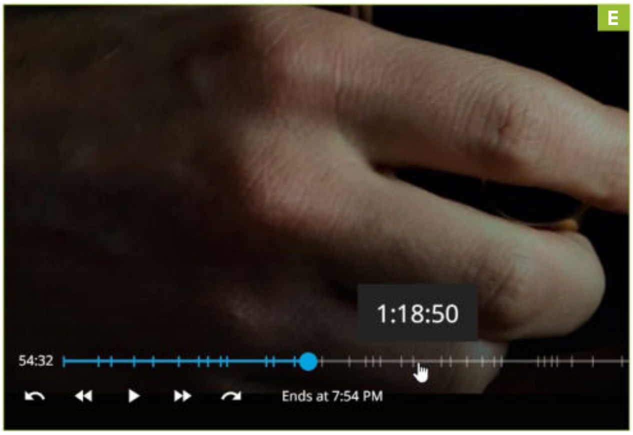

mercifully simple: roll your mouse over the offending title's thumbnail and you'll see a vertical ellipsis button appear in the bottom right corner. Click this to reveal a pop-up menu, as shown in [**Image D**] . Select Identify to bring up a search engine. Type the name of your title, plus the year of its release, and click Search to see if you can get the correct version. Select this, and click OK to update your library.

» Once you've identified a title, click it to access its own dedicated page, complete with synopsis, cast list, and access to any special features you've added (such as featurettes and deleted scenes)—see the annotation for details. If you've selected a TV show, you'll also be shown a list of seasons—click one to view a list of episodes from that season. You can play these directly from here, or click the i button next to it to access its own dedicated page.

» Returning to the main title, it's not just tags you can click on to reveal links to other titles in your collection—click on a person, for example, and you'll be shown a short bio, accompanied by convenient shortcuts to all their other credited appearances in items on your server, whether that's movies, TV, or music.

7

WATCH AND NAVIGATE

To view a title from here, click the play button. The title will quickly load and start playing—you'll briefly see playback controls on-screen, which will vanish until you move your mouse or click on your remote. You'll see the play/pause button is accompanied by additional jump controls allowing you to move between chapters or simply skip backwards and forwards 30 seconds at a time; alternatively, click on the timeline to jump to a specific part of the movie immediately. If you enabled TrickPlay for your library, you can drag the mouse back and forth along the timeline and use thumbnail images to help identify the spot where you want to go; if not, you'll have to make do with a time counter, as shown in [**Image E**] .

» You'll also see buttons on the bottom right, where you can adjust the volume, select a subtitle track and even switch audio track (click the music note button) if supported by the movie in question. Click the cog button to access even more options, including switching aspect ratio, playback speed, and playback quality.

8

KEEP EXPLORING

We've barely scratched the surface of what your Jellyfin server can do for you in part one of this tutorial. Next issue you'll be introduced to the finer points of configuring your server, plus you'll discover how you edit your libraries' rich text and image content using the powerful Metadata Manager. We'll also show you how to add support for streaming live TV from a suitable TV tuner, and you'll explore some of Jellyfin's brilliant plugins, including one that allows you to skip intros and credit sequences.

» To start exploring, click the hamburger button at the top of the screen. Select Dashboard where you can explore your server's settings, plus add additional libraries and set up dedicated user accounts for friends and family members. Start with a name, password, and what libraries they can access. Once created, you'll unlock a series of tabs that allow you to give them management capabilities, set parental controls and restrict access to specific devices, among other things. ⏻

REMOTE ACCESS 101

Want access to your Jellyfin libraries while on the road? Insecure connections are easy to configure—navigate to Dashboard > Networking, and verify that 'Allow remote connections to this server' is checked. Then, it's just a case of setting up a subdomain or dynamic domain that points to your home IP address, plus set up a port forwarding rule in your router to direct traffic on port 8096 to your Jellyfin server's IP address. When connecting from a client, use http://sub.domain.com:4096 to make contact.

If you have access to a proxy server like Nginx Proxy Manager (click 'Nginx proxy' under https://linuxformat.com/archives?listpdfs=1 for a comprehensive guide to setting it up), you can easily set up a secure connection without having to do any additional configuration in Jellyfin. From Nginx Proxy Manager, set up a simple proxy host redirecting traffic from your choice of subdomain to http (Scheme), PC's IP address (Forward Hostname/IP), 8096 (Forward Port). Flick both 'Block Common Exploits' and 'Websockets Support' switches to on, then switch to the SSL tab to assign a free LetsEncrypt SSL certificate to the redirect to avoid certificate errors.

Once in place, get access to your server by entering https://sub.domain.com into the Host box when connecting to a new server with your client.

Beat Games with Edge Game Assist

YOU'LL NEED THIS

MICROSOFT'S NEW GAME BAR WIDGET comes to the rescue of in-game web searchers

FOR MANY YEARS, almost as long as there's been an easily accessible internet and PCs could run two apps at once, gamers have been dumping themselves out of their games in order to search for info. Even in the era of cellphones that are themselves little supercomputers, people have preferred to use the desktop browser to look up cheat codes, boss weak spots, and more.

While games are generally more forgiving of you hitting the Windows key or Alt-Tab, it's still possible to get in a situation where the game refuses to restart after you switch back into it. This leaves gamers with a feeling of jeopardy—is the information you glean from switching to the browser valuable enough to risk losing your progress?

Microsoft has released an answer to this conundrum: a mini version of the Edge browser living in the Xbox Game Bar, an overlay you can summon using a keyboard shortcut. It's currently in Preview, which is kind of the equivalent of a late beta test, so you'll have to install it yourself, rather than it being an integral part of the system. Here's how. –**IAN EVENDEN**

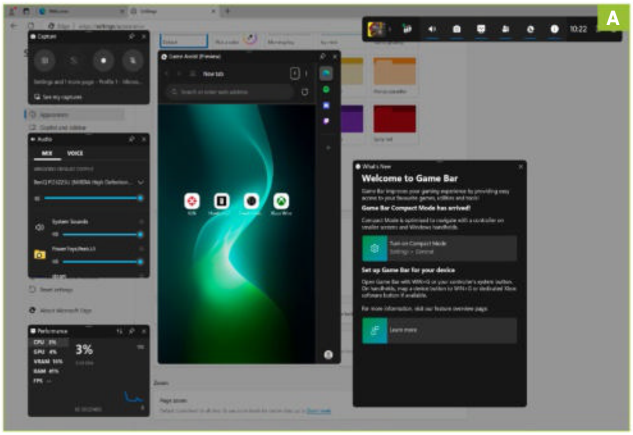
A
Welcome to Game Bar
Game Bar Compact Mode has arrived!
Set up Game Bar for your device
3%

1 THE WIDGET

You'll need Microsoft Edge version 132 or later for this, so make sure it's updated by opening the About Edge page from the browser's settings menu, and let it do its thing, restarting if necessary. From there, move up to Appearance in the list on the left-hand side, and click the Install button next to 'Game Assist (Preview)'.

2 THE BAR IS OPEN

You'll then be able to click 'Open', which brings up the Game Bar overlay with a browser window already open, A few game-related sites will already be on the shortcuts bar and services, such as Twitch and Spotify in the sidebar [Image A]. You don't need to do this every time, as you can use the Win-G shortcut or the Xbox button on a controller to open the Game Bar from within any game—or indeed, any other app—as the bar's not picky. If you have the Game Bar controller shortcut disabled in Windows Settings, you'll need to go to Settings > Gaming > Game Bar and slide the switch to enable it, otherwise you'll be stuck using the keyboard shortcut.

3 NEW GAME BAR FEATURES

There have been a few tweaks to the Game Bar itself recently, including a new compact mode [Image B] that's designed for handheld PCs running Windows, such as the ROG Ally or MSI Claw. You can find it by sliding on the toggle within Game Bar's settings menu, which is hiding behind the cogwheel icon in the main bar. Instead of having multiple windows open, including the new browser, the compact Game Bar limits you to one at a time, and provides larger icons so you can switch between them using touchscreen finger taps. Alternatively, you can move up and down between them using the left and right bumper buttons on a controller. The Home menu shows recently played games, along with shortcuts to Steam, Ubisoft Connect, and the Windows 11 Xbox app. All the other usuals are there, too, including performance metrics, Xbox friends, screen recording, and audio control.

» The Edge Game Assist browser appears in a larger window in compact mode, filling the screen from the start, rather than expecting you to use a mouse to drag it out to a larger size, which makes it the most user-friendly implementation of the Game Assist browser. Compact mode may be designed for portables, but it'll work just as well on your 32in 4K screen. The Xbox app itself also has a compact mode aimed at portable PCs—click the profile icon in the top left of the interface (you don't actually need to be logged in for this to work), and you'll find a Compact Mode toggle there. It doesn't do much other than collapse the sidebar down into icons, but it also makes tiles larger and easier to deal with on a touchscreen.

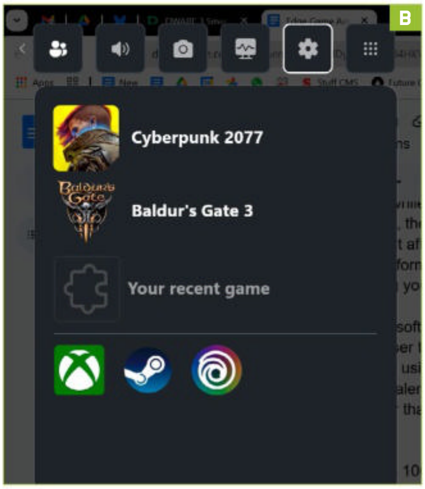

4 READY TO BROWSE

Back to the Game Bar. However you get into Edge Game Assist, what you're working with is a stripped-down version

© MICROSOFT

MEET THE GAME BAR

1. THE GAME BAR ITSELF
From here, you can open any of the Game Bar widgets, or use the Settings panel to switch into Compact mode.

2. CAPTURE
Capture screenshots or video of whatever you're doing, and choose whether or not to record audio, too.

3. AUDIO
Adjust the volume of speakers or the sensitivity of your microphone from here. Change the volume level of individual apps, too.

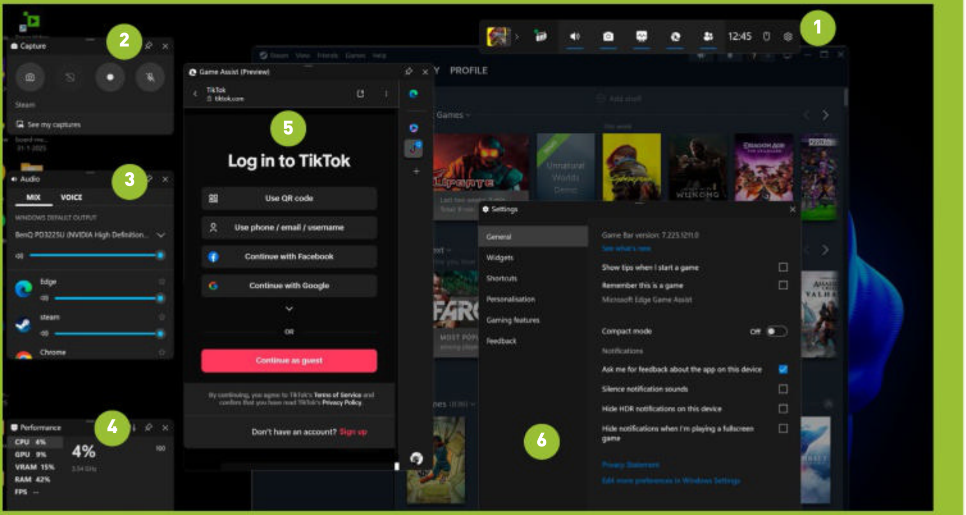

Log in to TikTok

- Use QR code
- Use phone / email / username
- Continue with Facebook
- Continue with Google

OR

Continue as guest

By continuing, you agree to TikTok's Terms of Service and confirm that you have read TikTok's Privacy Policy

Don't have an account? Sign up

CPU 4%
GPU 9%
VRAM 15%
RAM 42%
FPS --
4%

Game Bar version: 7.225.1211.0
See what's new
Show tips when I start a game ☐
Remember this is a game ☐
Microsoft Edge Game Assist
Compact mode Off
Notifications
Ask me for feedback about the app on this device ☑
Silence notification sounds ☐
Hide HDR notifications on this device ☐
Hide notifications when I'm playing a fullscreen game ☐
Privacy Statement
Edit more preferences in Windows Settings

4. PC PERFORMANCE
Ever wondered how hard your CPU and GPU are being worked while in a game? Pin this over whatever you're playing and watch the graphs.

5 EDGE GAME ASSIST
The stripped-back browser widget is found here. This can be used for undertaking a quick web search—an alternative to Alt + Tabbing to the main browser.

6 SETTINGS
From here, you can customize aspects of the user experience like altering the appearance of the Game Bar interface, setting your own shortcuts, and enabling HDR.

of the standard Edge browser, so you won't get things like the local weather report or MSN news feed that can clutter up the new tab page [**Image C**]. There's no option to swap it out for Chrome, Firefox, or any of the others, but modern Edge isn't that bad. It's certainly not like the dark days of later revisions of Internet Explorer, and as it uses the Chromium open-source project and the Blink rendering engine to build on, it's pretty fast, and certainly enough for a quick web search for potion recipes. This browser uses the same user profiles as full-fat Edge, which means it has all your history and browsing data, making it easier to access a site you've been to before, as it will probably autocomplete in the address bar.

5 IT REMEMBERS

The best thing about the Game Assist browser is that it remembers where you left off. When you go back to your game, the browser state is saved so that when you inevitably hit Win-G again, it picks up where it left off. This is helpful, because apart from the back and refresh buttons, there are few of the usual browser tools. There's no history menu, the three-dot menu contains nothing but the option to send feedback, and even opening a new tab requires a couple of clicks or taps (on the '1' icon that shows how many tabs are open, then New Tab). There's no right-click menu, so any hyperlinks you click will open in the same tab.

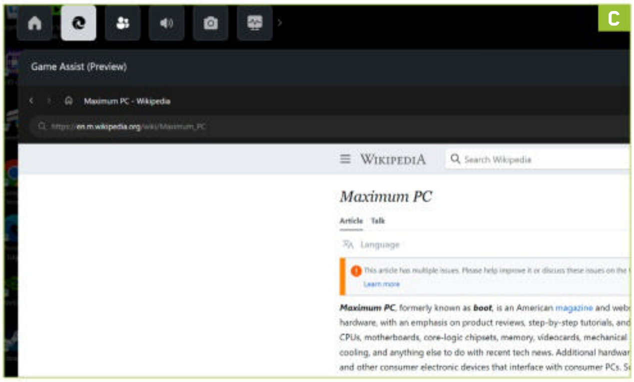

Maximum PC

This article has multiple issues. Please help improve it or discuss these issues on the
Learn more

Maximum PC, formerly known as **boot**, is an American magazine and web
hardware, with an emphasis on product reviews, step-by-step tutorials, and
CPUs, motherboards, core-logic chipsets, memory, videocards, mechanical
cooling, and anything else to do with recent tech news. Additional hardwar
and other consumer electronic devices that interface with consumer PCs. S

6 STRIPPED BACK

The stripped-back nature of the Game Assist browser point to its intended use. You don't want a resource hog popping up over the top of a game that depends on squeezing every last processor cycle out of your PC to achieve maximum performance, so there's little that can cause that. You can add web apps to the right-hand shortcuts sidebar (click the '+' icon and you get the likes of Outlook mail, AI image creation, Microsoft 365, Facebook Messenger, and lots more. These include games, so you can play while you play, but we found that even though it was possible to add the games to the sidebar, clicking on them didn't have any effect. We were able to access the Word documents stored in our OneDrive, however, as well as access TikTok, so this may just be a glitch that will be fixed with an update.

7 CLOSURE

As with all Game Bar widgets, all you have to do to close it is press Win-G again, or click on part of the screen that's not covered by the browser window. If you move to a Game Bar tab that's not the Edge one, and you're in Compact Mode, then the window will disappear, too. It's an efficient way to access the web while gaming without needing to rely on a cellphone or second screen, though the barebones nature of the browser doesn't make it ideal for regular use. ⏻

Pass on your passwords

YOU'LL NEED THIS

KEEPASSXC
(https://keepassxc.org)
Google Account

DEALING WITH THE DEATH of someone close is never easy. It's even harder when you share finances with them and/or need access to their social media to manage their memorial page.

On the face of it, the answer seems simple: just write down all your passwords and give them to a loved one in case something happens to you. However, it's not always that easy. As you add new accounts and change passwords, paper lists can quickly become outdated.

Certain platforms also support 'legacy' or 'trusted' contacts for situations like these. However, this is only a solution for accounts held by those services.

In this guide, you'll learn some more universal methods for passing on your credentials if something happens to you. You'll discover some best practices for writing down passwords, and learn which password managers support adding trusted contacts, as well as ways to do this yourself by using Google's Inactive Account Manager. **–NATE DRAKE**

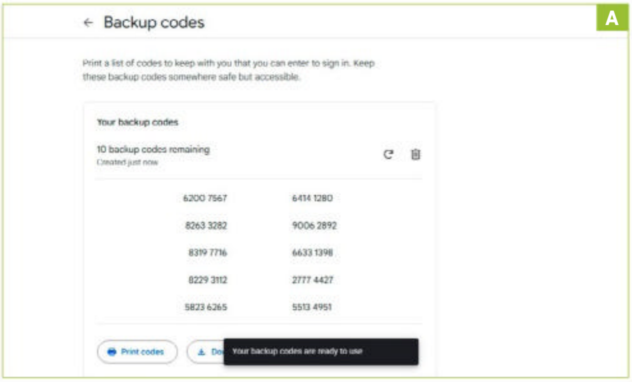

← Backup codes

Print a list of codes to keep with you that you can enter to sign in. Keep these backup codes somewhere safe but accessible.

Your backup codes

10 backup codes remaining
Created just now

6200 7567   6414 1280
8263 3282   9006 2892
8319 7716   6633 1398
8229 3112   2777 4427
5823 6265   5513 4951

Print codes
Do
Your backup codes are ready to use

1 WRITE IT DOWN

As keen as you may be to pass on your digital legacy, not every organization has a built-in way to do this.

» For example, Microsoft offers no way for you to designate a trusted person to log in to your Windows 11 devices. The same is true for passcodes for most mobile devices.

» Get started by making a list of all your devices and online accounts. When writing these down, follow the example of digital password managers by noting down the login URL, username (eg. your email address), and the password.

» If you log in to any accounts using 2FA (two-factor authentication), your legacy contact may still be able to sign in, provided they have the passcode to the mobile device on which the authenticator app is installed.

» Check if your provider supports 2FA backup codes. These can be used to access your account, even if your trusted person no longer has access to the authenticator app. The steps to do this will vary depending on your account type.

» For instance, for Google accounts, visit https://account.google.com. Next, choose 'Security' > 'Backup codes'. From here, you can generate 10 backup codes, then print them to add to your password list. [**Image A**]

» Google's security section also offers the option to add a secondary backup verification phone. Access this via https://myaccount.google.com/two-step-verification/phone-numbers.

» If you add another number, it will receive 2FA codes via SMS in the same way as your main number. Naturally, you should only do this with the number of someone you trust.

» If you use your browser to store passwords, check if these are stored online. For instance, Brave's 'Sync Chain' feature encrypts saved credentials using a 'recovery phrase'. This can be entered into another device to retrieve all your passwords. [**Image B**]

» Once you have a hard copy of your user credentials and 2FA codes (if necessary), place them inside a fireproof document wallet. Put them somewhere safe, such as alongside your Will.

2 USING A PASSWORD MANAGER

While writing down your credentials can be a good way to preserve your digital legacy, it can also be risky. If someone else finds the documents, they could compromise your online accounts.

» Unless you take elaborate precautions, paper records can also easily be destroyed. Recognizing this, certain manufacturers of cryptocurrency hardware wallets offer devices made of metal to store seed phrases and private keys. [**Image C**]

» Still, these can only hold a limited amount of information. If you already use a password manager, check with your provider whether they support legacy contacts.

» The steps taken to do this will vary depending on your provider. For example, NordPass has a 'Trusted Contacts' feature. To enable this, you need to log in to your account and enter the email address for your

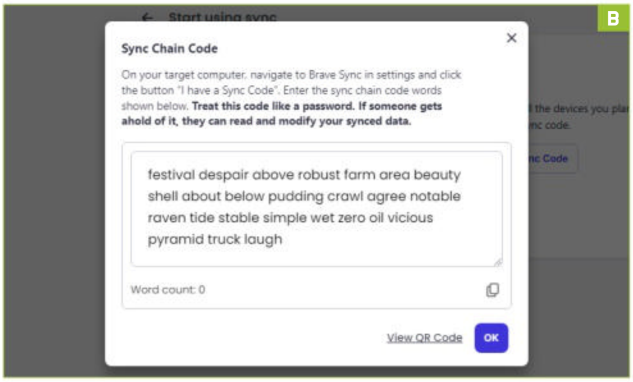

Sync Chain Code

On your target computer, navigate to Brave Sync in settings and click the button "I have a Sync Code". Enter the sync chain code words shown below. **Treat this code like a password. If someone gets ahold of it, they can read and modify your synced data.**

festival despair above robust farm area beauty shell about below pudding crawl agree notable raven tide stable simple wet zero oil vicious pyramid truck laugh

Word count: 0

View QR Code
OK

© KEEPASSXC.ORG

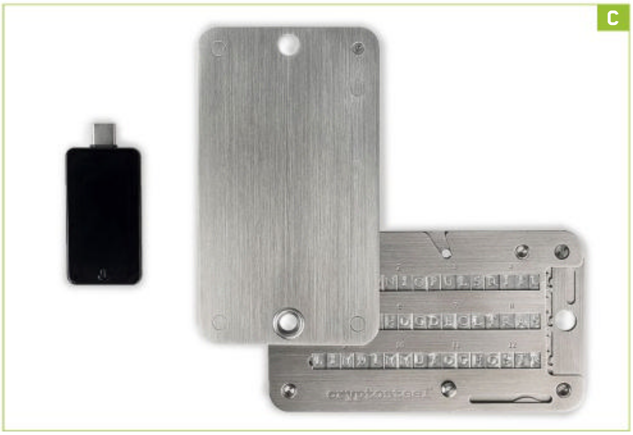

C

contact of choice. Unfortunately, this feature is available only to NordPass Premium subscribers. We found the same to be true for Bitwarden's 'Emergency Access' feature.

» If you don't want to pay, you can still provide your passwords to a trusted contact using the open-source manager 'KeePassXC', which is available via the Microsoft Store.

» To run it, you'll also need the latest Microsoft C++ Redistributable version (https://aka.ms/vs/17/release/vc_redist.x64.exe).

» Once both of these are installed, launch KeePassXC and choose 'Create Database'. Enter a meaningful database name, then choose 'Continue'. The default configuration options, ie. KDBX4 format with AES-256 encryption, are fine, but choose 'Advanced' to fine tune them, if you wish.

» When prompted to enter a password, click the dice icon to use KeePassXC's generator. Choose 'Passphrase' and note down the seven random dictionary words. Choose 'Apply Password' > 'Done'. [**Image D**] You can now pick where to save the password database.

» Before adding passwords, choose 'Groups' > 'New Groups' to create categories like 'Email accounts'. From here, you can choose 'New Entry', then add information, like your username and passwords. Use 'Notes' to record extra information, like 2FA backup codes. [**Image E**]

» The KeePassXC website has an excellent guide on getting started (https://keepassxc.org/docs/KeePassXC_UserGuide).

3 GETTING INACTIVE

Once you've added all your credentials to the KeepassXC database, hold Ctrl+F4 to close it.

» At this stage, you can store the database file securely offline. For instance, you could copy it to a USB stick. However, you're still risking this information falling into the wrong hands.

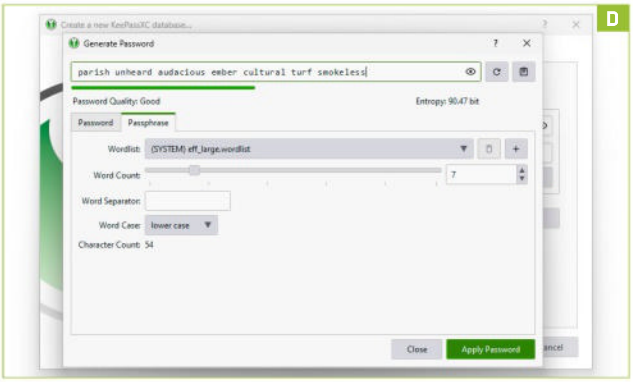

D

Create a new KeePassXC database...

Generate Password

`parish unheard audacious ember cultural turf smokeless`

Password Quality: Good

Entropy: 90.47 bit

Wordlist: (SYSTEM) eff_large.wordlist

Word Count: 7

Word Separator:

Word Case: lower case

Character Count: 54

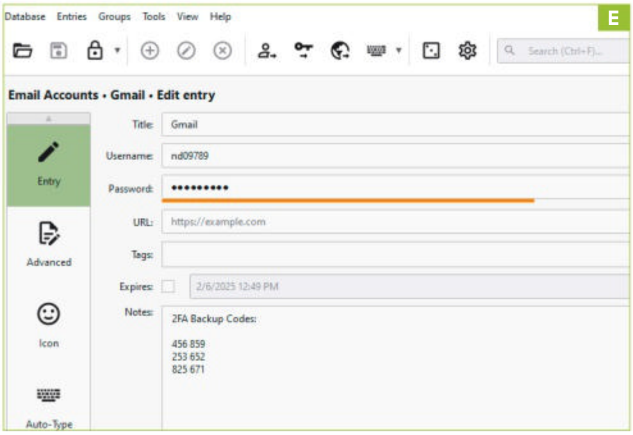

E

Email Accounts • Gmail • Edit entry

Title: Gmail

Username: nd09789

Password: ••••••••••

URL: https://example.com

Tags:

Expires: ☐ 2/6/2025 12:49 PM

Notes: 2FA Backup Codes:

456 859
253 652
825 671

» Another option is to use Google's 'Inactive Account' feature. This allows a trusted contact to be able to access the account if it's inactive for a certain period of time. You can use your existing Google account for this, or set up a dedicated one.

» First, go to https://drive.google.com, then drag and drop the KeePassXC password database you created earlier into the window. You may want to create a dedicated folder for it, eg. 'keepass-passwords', so that it's easy to find.

» Next, visit https://myaccount.google.com/inactive, and choose 'Start' to access the 'Inactive Account Manager'. Under 'Decide when Google should consider your Google Account inactive', find 'After 3 months of activity', and change the duration, if you wish. Periods of up to 18 months are supported.

» Select 'Next', then under 'Choose who to notify and what to share', click on 'Add Person'. Enter the email address of your trusted contact. You can now choose what aspects of your Google account to share, such as 'Chrome' or 'Contacts'. Select whatever you wish, but make sure that 'Drive' is highlighted. This ensures your contact can access the password database. [**Image F**]

» Google will prompt you to provide a phone number for your trusted contact. This is optional, but recommended. You can add a short message if you wish, eg. to inform them that they'll need KeePassXC to open the database. [**Image G**]

» Once this is done, make sure your trusted contact has a copy of your KeePassXC database password. The safest way to do this is to write it down and give it to them in person. Regularly update the database in your Google Drive so it reflects your current credentials. ⏻

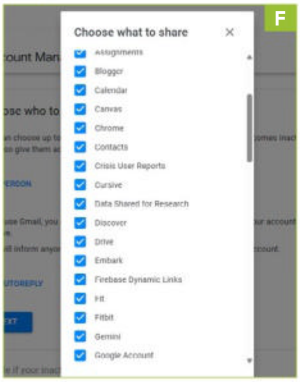

F

Choose what to share

☑ Assignments
☑ Blogger
☑ Calendar
☑ Canvas
☑ Chrome
☑ Contacts
☑ Crisis User Reports
☑ Cursive
☑ Data Shared for Research
☑ Discover
☑ Drive
☑ Embark
☑ Firebase Dynamic Links
☑ Fit
☑ Fitbit
☑ Gemini
☑ Google Account

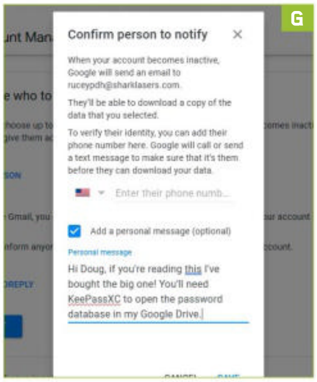

G

Confirm person to notify

When your account becomes inactive, Google will send an email to ruceypdh@sharklasers.com.

They'll be able to download a copy of the data that you selected.

To verify their identity, you can add their phone number here. Google will call or send a text message to make sure that it's them before they can download your data.

Enter their phone numb...

☑ Add a personal message (optional)

Personal message

Hi Doug, if you're reading this I've bought the big one! You'll need KeePassXC to open the password database in my Google Drive.

LAB NOTES

ZAK STOREY, CONTRIBUTING EDITOR

Samsung announces PCIe 5.0 SSD… finally

It's here. It's finally here

TO QUOTE STARCRAFT… hell, it's about time. It's finally here, ladies and gentlemen: the thing that Guy and I have been pining for literal months over many a Story of the Build video. We finally have Samsung's latest PCIe 5.0 x4 drive. Yep, at long last, nearly two years on from their initial release, Samsung has announced the arrival of its own SSD to join the club, and it comes with some top-tier specs.

We've got two samples in our office currently: the heatsinkless 4TB variant, and the heatsinkked 2TB variant, and boy do they punch. 14,800 MB/s read, 13,400 MB/s write, and random read/write IOPS at 2200K/2600K: it's a bit of a beast. Each drive comes with 1x 1GB of LPDDR4X cache per TB of storage, too, meaning the 8TB variant comes with a total of 8GB. Sammy's also using its own in-house controller (currently undisclosed, so we have no idea what's actually going on in there), combined with, perhaps unsurprisingly, its own TLC V NAND, appropriately dubbed V8.

Other top-line stats are fairly standard; you've got the usual bevy of data encryption feature sets, plus your bog-standard TBW endurance ratings for a drive of this size (1200 on the 2TB, and 2400 on the 4TB), and of course a five-year warranty. Pricing starts at $200 for the 1TB, topping out at $570 for the 4TB, with 8TB pricing news later in the year. It's that last drive that's particularly interesting, as it should be the first 5.0 8TB drive we've ever seen.

From what I can tell from first appearance, this is absolutely littered with tech. It's a lot more complex than some of the other 5.0 drives I've seen for traces and capacitors, that's for sure. Plus, both the NAND packages and controller are huge in contrast to the stuff Kioxia and Phison are producing. It's still in the lab undergoing stringent testing, but we'll let you know all about it next issue (you can bet your bottom dollar that it'll be going in a build, too).

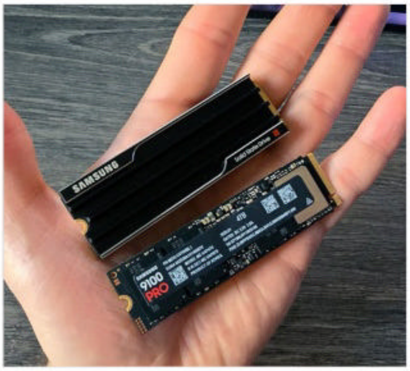

Heatsink, or heatsinkless drives available, up to 8TB, all of it sexy.

JEREMY LAIRD

Contributing Editor

Nvidia's new RTX 5070 graphics card has hit store shelves. Except, it virtually hasn't. Launch day has been a farce, with little availability and, as far as I saw, pretty much zero at the RTX 5070's supposed $549 MSRP.

As much as I don't like fake MSRPs, I understand them. You announce an attractive-looking MSRP. When the reality is a bit higher, and you're Nvidia facing limited competition, the assumption is that gamers will suck it up. Meanwhile, the competition has to offer lower MSRPs that they will be under pressure to stick to.

What I don't get is Nvidia 'launching' new RTX 50 GPUs with little to no supply. Some argue that's because all the volume at chip manufacturer TSMC is taken up by AI GPUs. I say that's total BS.

All RTX 50 GPUs are built on TSMC's N4 node. N4 is mature, has great yields, and is available in large capacity. Moreover, AMD makes its Zen 5 CPU dies on N4, and has no trouble supplying those.

Gaming graphics is small fry for Nvidia, such are the riches of its AI division. But that just means the cost of making sure there's a decent supply of new GPU models at launch is peanuts.

One day, Nvidia might not be quite so dominant in AI, and graphics cards will become more important again. If that happens, treating its gaming customers like crap could come back to bite them.

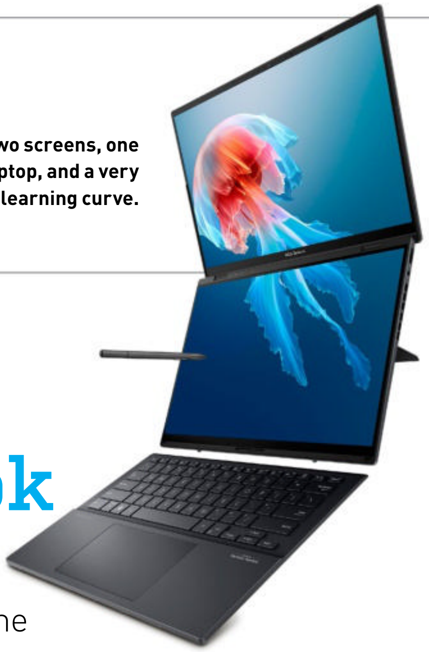

Two screens, one laptop, and a very steep learning curve.

Editor's Pick: Asus Zenbook Duo 2025

You won't want to go back to one

ONE DEVICE, two very different aspects. That's the Asus Zenbook Duo 2025—a tale of two halves. In simple terms, it's a dual-screen laptop with a clip-on keyboard and great hardware.

That was true of the previous 2024 model. This revised-for-2025 device largely carries over the same beautifully engineered chassis containing those dual 14-inch OLED touchscreens. You also get a super-slick wireless clip-on keyboard.

The Asus Zenbook Duo 2025 looks like a pretty conventional laptop, albeit slightly portly. At over 3.5 lbs, it's also heavier than many 15-inch systems, such as a MacBook Air 15. But the keystrokes feel better than you'd expect for a clip-in board, and it works wirelessly when you detach it, enabling all manner of intriguing ergonomic setups.

What's more, the alloy chassis has a sturdy-feeling hinge and a kickstand that allows you to prop both screens up vertically. That's fabulous for some use cases—for instance, running collaboration software on the top screen, and your documents and apps on the bottom screen. Going back to a single-screen laptop once you get used to the Zenbook Duo feels awfully constraining, that's for sure.

Of course, those dual-OLED panels shine. The screens are identical 14-inch OLED panels with 2,880 by 1,800 pixels each, a peak HDR brightness of 500 nits, and eye-popping visuals. They also run at up to 120Hz, which is important not just for things like scrolling web pages and documents, but for ensuring that the touch input with the bundled Asus Zenpen 2 is responsive, precise, and lag free.

As for performance, cramming the latest Intel Arrow Lake CPU under an OLED screen was never going to be a recipe for peak performance. But the top-spec Intel Core Ultra 9 285H is a powerful CPU, though as an Arrow Lake model as opposed to Lunar Lake, it has a relatively weak NPU that doesn't qualify for Windows Copilot+ AI assistant functionality.

Then there's software—the other half of the Zenbook Duo equation. Windows 11 has never been well-optimized for touch. This laptop only serves to highlight that shortcoming. In theory, that's meant to be offset by Asus's suite of touch-optimized apps and features, all designed to make the most of the dual-screen hardware.

You can see what Asus is trying to do, including a highly modular and user-configurable touch control interface that in theory can be set up to suit almost any preference. You even get ready-made configurations for specific apps.

However, in practice, the learning curve is steep, and even once you've scaled that peak, the results are a little patchy and buggy. A slick, fully polished touch experience, this is not. There's a slight vibe of concept hardware or a design study, despite this 2025 model being a second-generation product.

Some of that is thanks to Windows itself being half baked, and some of it is surely down to Asus. In the end, it doesn't matter; the result is at least occasional frustration. That doesn't make the Asus Zenbook Duo 2025 unbuyable, but for a machine this expensive, it does make for some significant caveats.

If you're willing to put up with some pain, there are rewards to be had. This system can do things that a conventional laptop can't dream of. But in return, you'll have to accept aspects that feel clunky and unfinished. **–JL**

$1,699, www.asus.com

© ASUS, 2K GAMES

Reviewed...

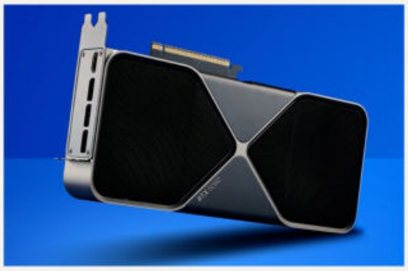

74 Nvidia GeForce RTX 5080

76 Dough Spectrum Black 32 OLED

78 Gigabyte Z890 Aorus Elite WiFi7 Ice
80 HP Omen 35L
82 Dell XPS 13 9350
84 Asus Zenbook A14
86 Gigabyte X870 Aorus Elite WiFi7 Ice
88 Logitech G RS Wheel Hub & Track Wheel
89 NZXT Lift Elite Wireless
90 Civilization 7

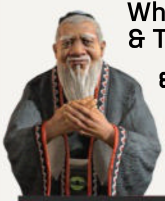

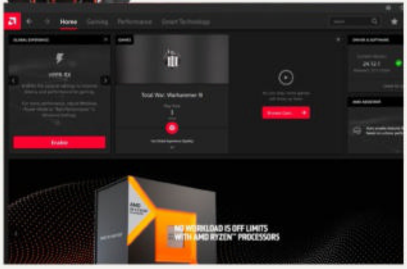

92 AMD Software: Adrenalin Edition vs Intel Graphics Software

Nvidia GeForce RTX 5080

Built for the Multi Frame Gen gaming experience

HALF THE MONEY for half the GPU. We give you the Nvidia GeForce RTX 5080, the second coming of Nvidia's Blackwell series, and by most measures an RTX 5090 sliced in half. It's also uncomfortably similar to the last-gen RTX 4080 Super.

You get just five percent more cores than the RTX 4080 Super, with 10,752 in total. Despite rocking the same TSMC custom N4 lithography, it's actually the smaller chip, if only by a smidge. There are 45.6 billion transistors inside GB203, where there are 45.9 billion inside the older AD103 chip.

Meanwhile, it's the same 16GB of VRAM and 256-bit memory bus here as the RTX 4080 cards, except this time it's GDDR7 instead of GDDR6X running at 30Gbps versus at best 23Gbps for the previous Ada cards. So there's a fair chunk more memory bandwidth available, at least.

There are some other tweaks inside the GB203 silicon that separates it from the AD103 chip of the RTX 4080 Super. There are more texture units, which means more texture processing power, and more L1 cache, although you are looking at the same 65MB of L2 cache across the chip. Nvidia is throwing a bit more power at the card, too, with the TGP rated at 360W versus 320W for the RTX 4080 Super. That means the recommended PSU specs are up by 100W.

The internal architecture is essentially the same as the RTX 5090 we reviewed last issue. Suffice to say, the big change is the fact that the shaders are now to be given direct access to the Tensor Cores of an Nvidia GPU—rather than relying on CUDA programming—which will allow a level of AI game integration we've not seen before.

Then, of course, there's Multi Frame Generation. Entirely locked down to the RTX 50-series for now, the skinny is that MFG is only possible because of the power of the 5th Gen Tensor Cores and a few other architectural revisions of the Blackwell family. The ability to AI-generate up to three extra frames between each pair of full 3D pipeline render frames is impressive on its own, but being able to do so without adding a ton of extra latency into the picture and pacing it all perfectly is the real magic.

MFG FTW

Ultimately, MFG is the 5080's defining feature. Without it, the 5080 would be very little advance over the old 4080 boards. Specifically, the overall gen-on-gen performance boost is a reliable 15 percent at 4K across our test suite. That's barely enough to feel in-game. If you were hoping for RTX 4090 rendering performance from the second-tier RTX Blackwell card, you're going to be disappointed—unless you own a 40-series GPU like a 4080, in which case you'll be pleased to find that desire for an expensive upgrade just evaporated.

What's more, the performance delta versus the RTX 4080—as with the RTX 5090—shrinks as we drop down the resolution scale. At 1080p and 1440p it drops to nine percent and a touch under 14 percent respectively. The apparent performance picture changes dramatically with Multi Frame Generation, especially in combination with upscaling. Going from 20fps at 4K native to 130fps with RT Overdrive in *Cyberpunk 2077* and DLSS Quality with 4x MFG really does give you the generational improvement we've been craving.

The latency is reasonable, too, at 67ms. But it's higher than what you'll get from 130fps when fully rendered, and you'll have to put up with some minor image flaws that are introduced by all that AI processing. Even Nvidia's new transformer model for upscaling isn't perfect, and Frame Gen creates its own visual niggles, too. Moreover, you do need a decent base frame rate with tolerable latency before adding Frame Gen. As this GPU ages, that could become a problem, especially at 4K.

So, the work Nvidia has done in making Multi Frame Generation work is truly impressive, but the question is whether it's enough. The raw rendering performance of the 5080 is a very small step over the last generation, and no matter how many more MFG frames you add, they aren't the same as fully rendered frames. Time will tell how much of a problem that proves to be. **–DAVE JAMES AND JEREMY LAIRD**

VERDICT **7**

Nvidia GeForce RTX 5080

⊞ **GENERATION GAME** Multi Frame Gen can be magical; Not as expensive as feared.

⊟ **GAMING THE NUMBERS** Frame gen isn't always the answer; Raw rendering grunt disappoints.

$999, www.nvidia.com

BENCHMARKS

	Nvidia Geforce RTX 5080	Nvidia Geforce RTX 5090	Nvidia Geforce RTX 4090	Nvidia Geforce RTX 4080 Super
Black Myth Wukong 4K High (fps)	55	**82**	64	47
CybeprPunk 2077 4K RT Ultra (fps)	36	**58**	45	31
F1 24 4K RT Ultra (fps)	65	**110**	81	56
Homeworld 3 4K Epic (fps)	98	**111**	102	89
Talos Principle 2 4K Ultra (fps)	53	**79**	63	47
Warhammer III 4K Ultra (fps)	66	**101**	72	54

Best scores are in bold.

SPECIFICATIONS

CUDA cores	10,752
Boost clock (GHz)	2.62
Tensor core TOPS	1,801
Ray tracing core TFLOPS	171
Memory	16 GB GDDR7
Memory bus width	256-bit
Total Graphics Power (watts)	360
Power connector	1x 450 W PCIe Gen 5 OR 3x PCIe 8-pin adapter
Price	$999

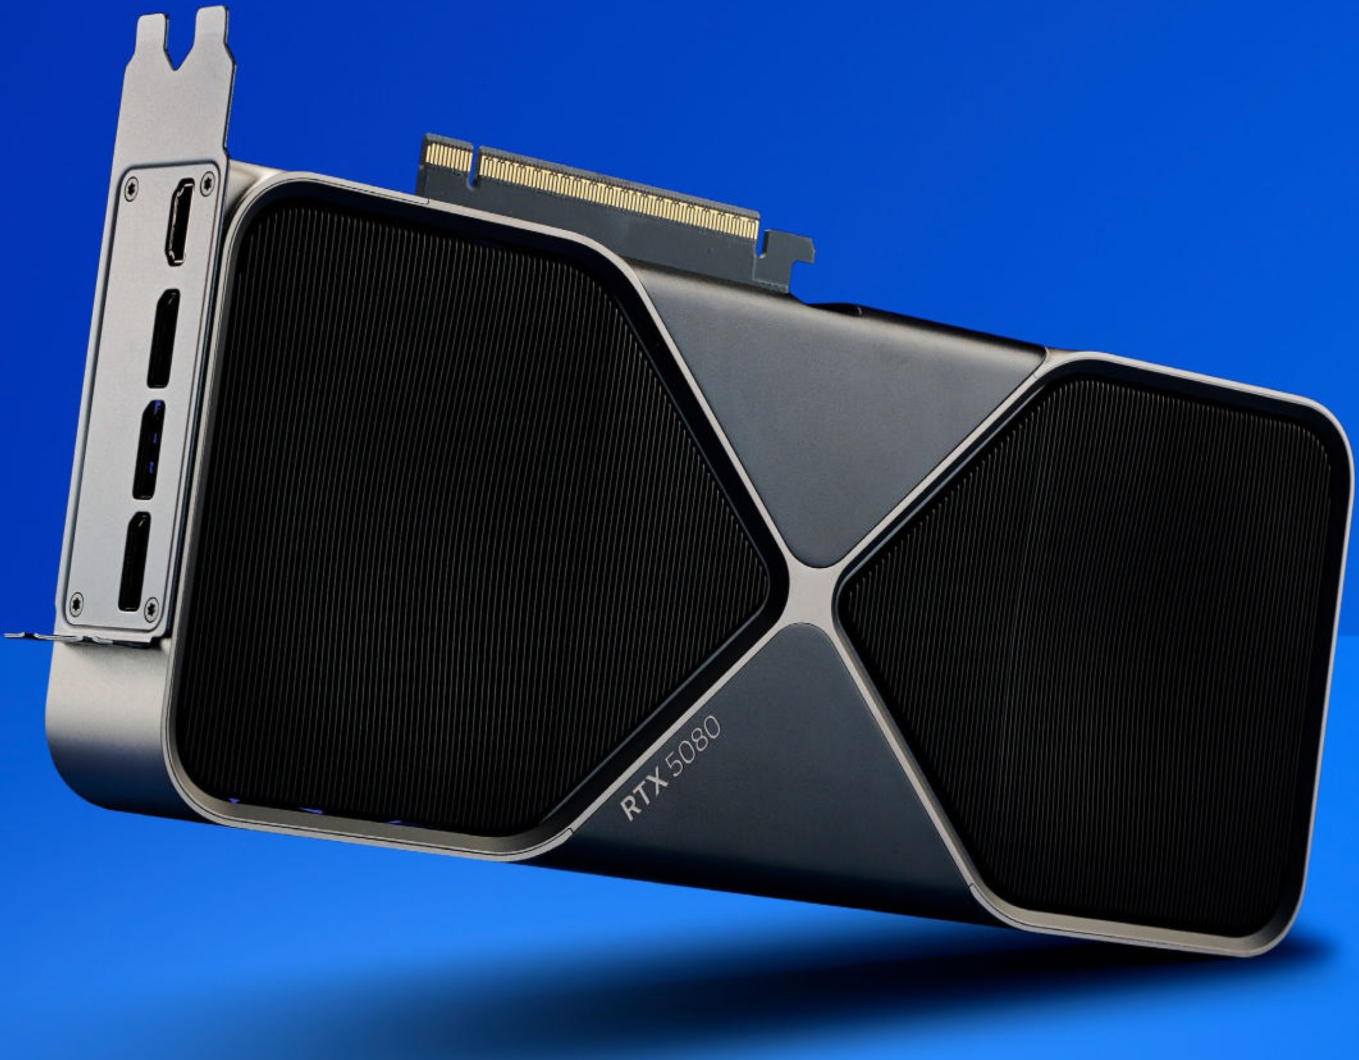

The RTX 5080 looks a lot like the 5090, but it's far less performant.

Dough Spectrum Black 32 OLED

OLED gets Gorilla Glass treatment

IT'S HARD TO KNOW what to make of Dough. You want to like it as an independent start-up. But under its former name of Eve, it — or rather its—customers, suffered from some widely publicized stumbles.

Whatever the company's past, with the Dough Spectrum Black 32 OLED, you sense care and investment in terms of design and attention to things like the aesthetic and build quality that big brands just don't give you. This 32-inch 4K OLED monitor sports a lush metal rear chassis and a gorgeous, beautifully engineered metal stand (an optional $100 extra) where monitors from mainstream brands tend to feel a bit cheap and plasticky, even at this premium end of the market.

It also delivers the numbers, specs-wise. Dough has chosen the LG WOLED panel option, as opposed to Samsung's QD-OLED. There are pros and cons to both, but the LG panel in 4K format is the company's 3rd generation WOLED technology, and closes the gap to Samsung for full-screen brightness.

If anything, it's a bit brighter full-screen for white tones, though QD-OLED probably has the edge for color brightness. That's because WOLED, as the name implies, has an additional white subpixel to boost brightness that does its best work with bright tones, which QD-OLED gets boosted by quantum dot tech that applies across the color spectrum.

Long story short, this monitor's 240Hz refresh and 275 nit full-screen and 1,000 nit peak HDR (in a three percent window) brightness ratings are in line with other 3rd Gen WOLED monitors, something that also applies to the speedy 0.03ms pixel response performance.

As for connectivity, you get HDMI 2.1 and DisplayPort 1.4, but no USB—at least, not with this model, which goes for $1,099, or $1,199 with the stand. A 'Hub' model, which adds that functionality and a few other features, is coming, but will be $1,299 without the stand.

But what about performance? The sRGB calibration looks great, with accurate colors and plenty of punch. There's a little ABL-driven brightness variation, depending on how much of the screen is ignited. But it's a minor issue.

The Gorilla Glass panel cover is likewise sweet, enabling a sense of heightened contrast, without excessive reflectivity. In fact, it makes for the best 'glossy' implementation of an OLED we've seen. You would have to use this monitor in bright ambient conditions to choose the matte version Dough will also be offering. The pixel density is also fab.

However, the problems come when you toggle HDR. While I updated to the latest firmware that supposedly fixes an HDR brightness issue, my review sample still lacked punch in HDR mode. What's more, the SDR tone mapping in HDR mode is a bit of a mess. The upshot is that you have to switch between SDR and HDR modes depending on content type.

I could live with that if it weren't for the price and lack of USB-C. If it were $999 with a stand and USB-C, this would be my favorite 4K OLED. The design and build is lovely, and the LG WOLED panel is fabulous. Fix the SDR tone mapping, lower the price, and add a USB-C interface, and this will be a killer product. As it is, there are just a few too many caveats for a firm buy recommendation. **–JEREMY LAIRD**

VERDICT

7

Dough Spectrum Black 32 OLED

➕ **GLOSSY** Lovely build and design; WOLED 4K panel; Gorilla glass is great.

➖ **LOSSY** Calibration needs a little work; No USB-C connectivity.

$1,199 www.dough.tech

SPECIFICATIONS

Screen size	32-inch
Panel type	LG WOLED
Resolution	3,840 x 2,160
Brightness	275 nits full screen, 1,000 nits max HDR
Response time	0.03 ms
Refresh rate	240Hz
HDR	DisplayHDR 400 True Black
Features	LG WOLED panel, Adaptive Sync, 1x DisplayPort 1.4, 2x HDMI 2.1, Gorilla Glass 3

© UNSPLASH

The Dough's Gorilla Glass panel cover enables a sense of heightened contrast, but without excessive reflectivity.

Gigabyte Z890 Aorus Elite WiFi7 Ice

A slick board, but the AMD alternative looms large

IT HAS BEEN a somewhat bizarre launch for Intel, with its Core Ultra desktop processors crashing and burning (thankfully not physically this time) in reviews, and sales figures are said to be poor. Indeed, it's only because of ongoing microcode updates and fixes for Intel's new Arrow Lake chips that we'd even consider opting for Intel's new platform and a motherboard such as the new Gigabyte Z890 Aorus Elite WiFi7 Ice.

BIOS updates, driver updates, Windows updates—they've been coming thick and fast for Z890 motherboards in the hope of improving performance. It's time to see if the latest tweaks for the platform—specifically the Gigabyte Z890 Aorus Elite WiFi7 Ice—make it worth considering, particularly when compared with its AMD X870 counterpart that costs a touch more and the review for which you can catch a few pages later this month.

Like that AMD board, the focus is on Wi-Fi 7, PCIe Gen 5 support, tool-free features, and high-speed USB Type-C ports. Being Intel, that means Thunderbolt 4. Oh, and that signature all-white icy aesthetic, which even extends to the PCIe and memory slots, as well as small details on the PCB. It's pretty impressive, considering this board leaves you with change from $300.

That said, despite having been released recently and carrying Intel's flagship chipset, you only get one PCIe 5.0 M.2 port, while the AMD equivalent has three, even if only one of those won't steal lanes from your graphics card. Both motherboards have four SATA ports, which is probably plenty, given that interface represents a dying breed.

Gigabyte's tool-free features are everywhere, from the push-button graphics card release of PCIe EZ-Latch Plus to its equivalent on the M.2 ports, heatsinks, and a quick-connecting Wi-Fi antenna cutting the time it takes to get a rig built and jump into your favorite game.

Moreover, features such as a POST code display and power and reset buttons are usually reserved for pricier boards, but, as with the AMD sister board, Gigabyte has included them. The I/O panel has nine Type-A USB ports, although there's only one Type-C Thunderbolt 4 port, while most USB4-equipped AMD boards have two of those.

Gigabyte's EFI is also included in the white theme. It's not as modern as Asus or MSI's updated efforts, but it's still easy to find useful settings. As for performance, operating temperatures running our Core Ultra 9 285K are down a couple of degrees versus the Asus ROG Maximus Z890 Hero and MSI MEG Z890 Ace at a peak of 50 C, though the peak PCIe 5.0 SSD temperature was warmer than the Asus board at 72 C, but below the 77 C recorded by the MSI MEG Z890 Ace.

Gaming performance is higher overall, and there are also smaller uplifts in Cinebench, Blender, and 7-zip, each time taking the performance crown. The combination of BIOS updates and the latest Windows updates all contribute, and there are more fixes to come. Power consumption was largely in line with other Z890 boards we've tested.

It's been tricky to recommend anything Z890-related, but the benchmarks here show that the situation for Arrow Lake has improved. However, if you're building a gaming PC, there's only one CPU you should be pairing a $300 motherboard with: AMD's Ryzen 7 9800X3D.

Nonetheless, Intel still offers a stong all-round computing proposition, and Gigabyte has taken the white theme to the extreme, plus shoehorned in way more features than we expected, given there's actually change from $300. Its AMD equivalent does offer a couple of extra features, but if $300 is your limit, and you want a white Z890 motherboard, this is easily the best option. **–ANTONY LEATHER**

BENCHMARKS

	Gigabyte Z890 Aorus Elite WiFi7 Ice	MSI MEG Z890 Ace	Asus ROG Maximus Z890 Hero
Cyberpunk 2077 (fps)	**144**	104	103
Baldur's Gate 3 (fps)	**106**	94	94
Gaming power draw (W)	91	94	**83**
Cinebench R24 multi-thread (pts)	**2,466**	2,434	2,390
Blender 4.2.0 (samples per minute)	**173**	169	169
Peak VRM temperature (degrees C)	**50**	58	52
Peak SSD temp (degrees C)	72	77	**66**

Best scores are in bold. All boards tested with an Intel Core Ultra 9 285K CPU and 32GB Corsair Vengeance Pro DDR5-6000

VERDICT

8

Gigabyte X870 Aorus Elite WiFi7 Ice

+ ICE IS NICE Great design; Good feature set; Bug fixes have improved Arrow Lake.

− STAY FROSTY Do you really want to go Intel?; Just one PCIe 5.0 SSD slot.

$270 www.gigabyte.com

SPECIFICATIONS

Socket	LGA1851
Chipset	Intel Z890
CPU compatibility	Intel Core Ultra 200 desktop
Form factor	ATX
Memory support	DDR5-4800 to DDR5-9200+(OC), up to 256 GB
Storage	4x M.2, 4x SATA
USB (rear)	1x Thunderbolt 4 Type-C 40 Gbps, 2x USB 3.1 Type-A 10 Gbps, 3x USB 3.0 Type-A 5 Gbps, 4x USB 2.0
Display	1 x Thunderbolt 4, 1x DisplayPort
Networking	Realtek 2.5G LAN, Wi-Fi 7
Audio	Realtek ALC1220

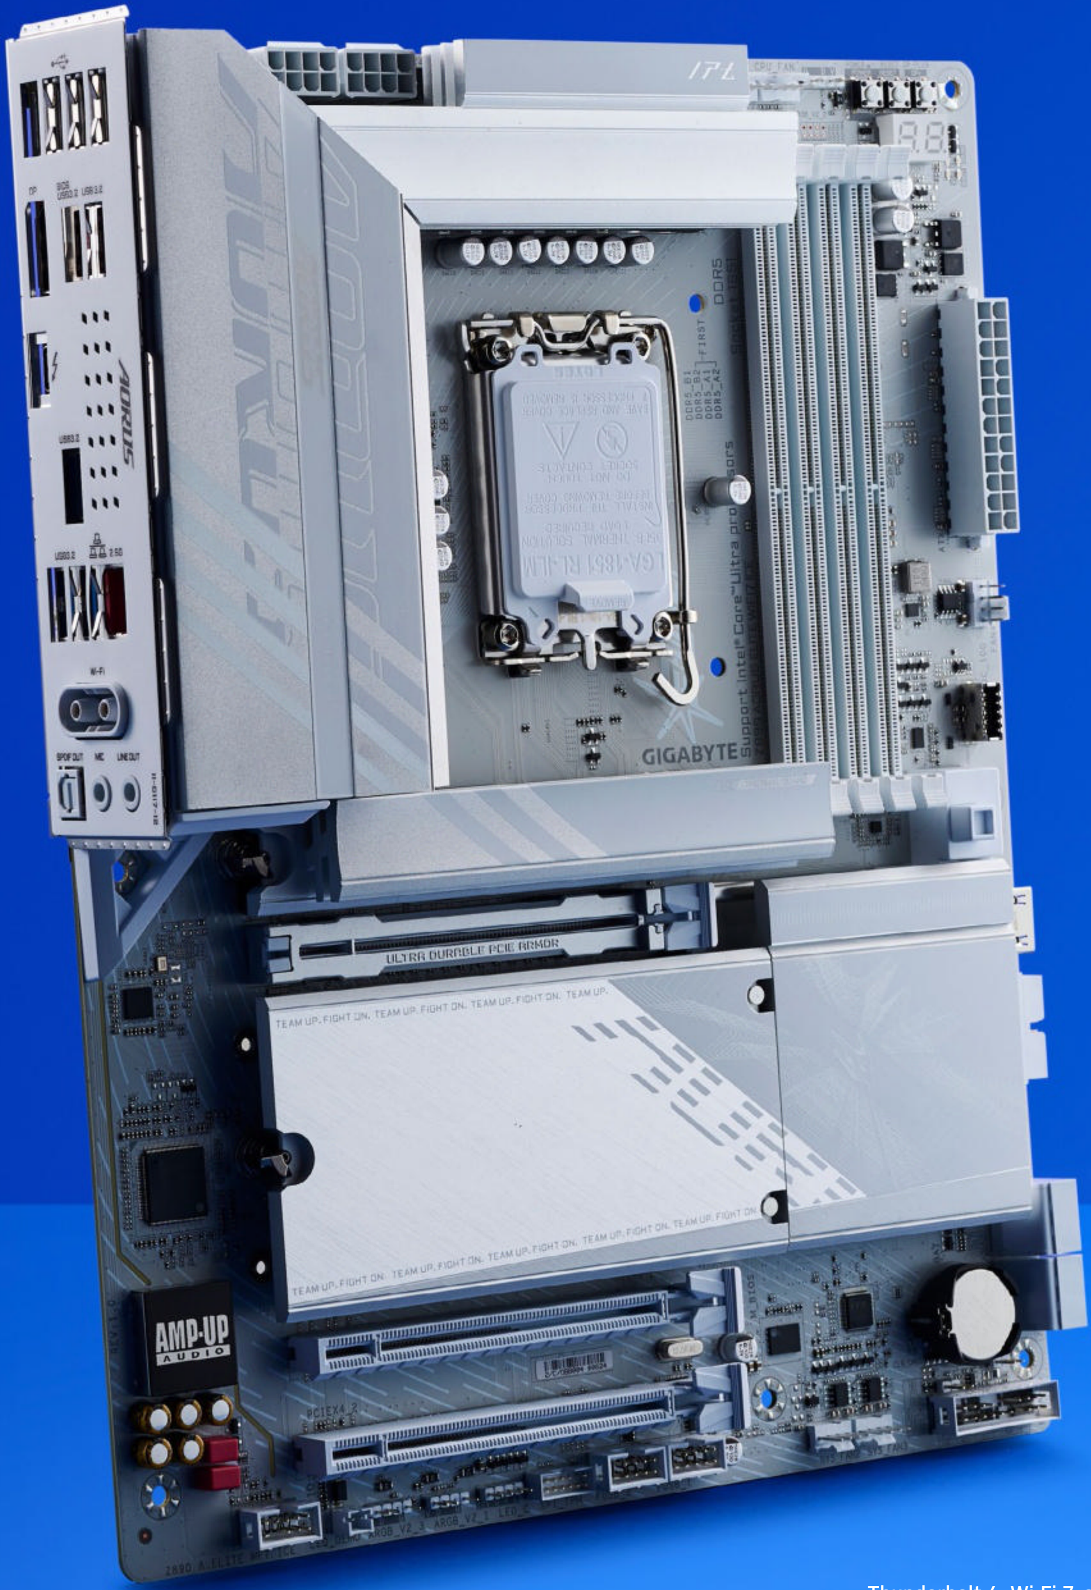

Thunderbolt 4, Wi-Fi 7, and PCIe 5.0 all feature, as do tool-free features and an all-white design that even extends to the EFI.

HP Omen 35L

HP's Omen 35L is a gaming PC with a love of tidiness

HP'S LATEST OMEN is a desktop tower that contains some serious power, yet unlike many other powerful gaming PCs, it doesn't take up too much space. It's not exactly compact—we're not talking Micro ITX levels of miniaturization here—but it doesn't feel large, even when it's on your desk. Add to this its remarkably quiet AIO loop cooler and GPU fans, and HP has delivered a powerful PC that's particularly easy to live with.

The RTX 4080 Super that sits in the Omen 35L's PCIe x16 slot is not known for being quiet, though it does depend on the manufacturer. The card is enclosed in a metal cage that matches the white coloration of the casing—and everything else apart from the motherboard—so it's hard to tell who made it, but it appears in CPUs as one of HP's own design. If that's true, the venerable PC maker has done a good job on the cooling system, as while it's certainly not silent, it produces a gentle rushing sound even when under load that's won't take your attention away from what's happening on the screen.

There's more subtlety on display elsewhere on the Omen. While the front fans are RGB, they're kept behind a screen so that their colors show through in a diffuse way. There are Omen logos, but they're kept to wordmarks, apart from the diamond logo on the CPU cooler. It's a classy way to brand your gaming PC, and the colored lighting from all angles lights up the white case and components of our review model in a really classy way. The lights go out when the PC has been idle for a while, too, popping back into life when you resume work in a pleasing way.

The RTX 4080 can produce over 50fps in *Cyberpunk 2077* at 4K and in Ray-tracing Overdrive mode—you could probably get that to 60 if you switched DLSS from Quality to Performance mode, but what's really interesting here is the CPU. It's a second-generation Core Ultra for the desktop, AKA Arrow Lake, and this 20-core (eight P, 12 E, turbo up to 5.5GHz) chip produces some appealing results. In Geekbench 6 testing, it posted a multi-core score that was greater than the mighty AMD Ryzen 7 9800X3D, leaving various Xeons, Threadrippers, and even the Apple M1 Ultra in its wake. It produced identical single-core results to AMD's latest 3D chip, and beat the Core i9-13900K. We look forward to seeing what Ultra 9 chips from this generation can do.

Elsewhere, the Omen 35L is nicely specced. There's Wi-Fi 7, 2.5G Ethernet, a single Thunderbolt 4 port (at the back), and easily accessible USB 3.2 Type-C and Type-A ports connectors at the front. The only thing holding it back—and it's entirely possible that this can be fixed by a configuration tweak or just wiggling it around until it's seated properly—is the SSD. Despite advertising itself as a PCIe 4.0 model, it produced PCIe 3.0 speeds in our tests, topping out at 300MB/s. This doesn't make a huge difference when gaming, though a faster drive may decrease loading times slightly, but there's a clear upgrade to be made here, along with slipping a GeForce RTX 5000-series card in to replace the 4080.

Those are the kind of changes you'd make after owning it for a couple years, however, and on face value, the Omen 35L is a speedy gaming PC presented in a very attractive way. What more can we ask for? **–IAN EVENDEN**

VERDICT
8

HP Omen 35L

➕ **GOOD OMEN** Good looks and performance in a reasonably sized case.

➖ **BAD OMEN** SSD performs less well than expected; Still quite expensive.

$2,900, www.hp.com/us-en/shop

SPECIFICATIONS

CPU	Intel Core Ultra 7 265K
GPU	Nvidia GeForce RTX 4080 Super 16GB
RAM	32GB DDR5-6000
Storage	2TB PCIe 4.0 SSD
Networking	Wi-Fi 7, Bluetooth 5.4, 2.5Gb Ethernet
Front panel	1x USB 3.2 Type-C 10Gbps, 2x USB 3.2 Type-A 5Gbps, 1x 3.5mm audio
Rear I/O	1x Thunderbolt 4, 2x USB 3.2 Type-C 10Gbps, 2x USB 3.2 Type-A 5Gbps, 4x USB 2.0, Ethernet, audio
Dimensions	8.27 x 16.06 x 16.14 in
Weight	32lbs

OMEN
35L
GEFORCE RTX
intel CORE ULTRA 7

The Omen 35L is neatly laid out inside, with most cables tucked out of the way.

Is Intel's Lunar Lake chip enough to make you forget about the Qualcomm Snapdragon X Elite?

Dell XPS 13 9350

OLED option gives battery life a beating

THE DELL XPS 13 is an old favorite. While its latest minimalist design is something of a divider in terms of aesthetics and keyboard, our high expectations remain. That's doubly true thanks to its new Intel Core Ultra 7 258V CPU, otherwise known as Lunar Lake. With the recent Snapdragon X Elite-powered Dell XPS 13 (9345) breaking records for battery life, the question is whether Intel can take the fight to Qualcomm. Let's find out.

Dell has kept the InfinityEdge bezels and flat keyboard deck of the previous few XPS models for the new Intel Core Ultra 200V-powered XPS 13. That's either excellent or disappointing, depending on how you've felt about the last couple of generations of XPS 13 models.

While this is a sleek look, it does lose some of the functionality of discrete function row keys and a visible border between the touchpad and the keyboard deck. Like the other XPS laptops with this keyboard design, the lack of space between keys and click activation makes for a difficult typing experience.

The borderless touchpad measures roughly 3.9 by 2.9 inches—smaller than most modern laptops. While there is a subtle difference between the touchpad and the keyboard deck, it can be difficult to determine the edges, especially if you're used to scrolling by feel. The touchpad haptics are helpful for this, but unfortunately, the function row still has no haptics.

Still, the minimalist design of the XPS 13 is elevated by just how light and thin the laptop is, making it feel more futuristic than ever. Coming in Platinum (white) and Graphite (black), the recent XPS design feels more iconic with every iteration. I wish we'd get a third colorway option, but I can't say I'm disappointed with the two we have.

Connectivity isn't the XPS 13 9350's strongest point, featuring just two Thunderbolt 4 USB Type-C ports with DisplayPort 2.1 and power delivery. If you rely on any USB-powered accessories during the day, you'll probably want to invest in a high quality USB Type-C hub or laptop docking station to suit. But if you want a highlight, that'll be the OLED display.

The 13.4-inch Tandem OLED display clocks in at 2,880 by 1,800 pixels and 60Hz refresh. More importantly, we tested it at an astounding 144.6 percent of the DCI-P3 color gamut with a Delta-E accuracy of 0.28. Those are figures that will blow away even professional grade desktop LCD monitors. OLED really is becoming a fantastic display technology. That said, brightness comes in at 377 nits—decent, but not exactly eye-popping.

As for the side-firing Realtek quad-speaker array, it cranks out decent volume. But the audio quality is distorted, as if you're listening through a tin can. It's disappointing for a premium laptop, and not nearly as good as, say, Apple's entry-level MacBook Air's audio. If you're using the XPS 13 for a lot of video or audio streaming, you'll want to use headphones or some decent external speakers.

Performance-wise, the new Intel Core Ultra 7 258V processor isn't the fastest, but in my experience, it's powerful enough for multiple programs, over 20 Chrome browser tabs, games downloading in the background, and some light Photoshopping on top of that. More importantly, it makes for staggering battery life—at least with an LCD display.

We notched up over 18 hours of web browsing with the 1080p LCD model. However, this OLED variant doesn't last half as long, expiring in just under nine hours. That's a serious price to pay, and creates a quandary. In our experience, OLED does normally hit battery life a bit, but not this badly, so perhaps the XPS's 'Tandem' OLED tech is to blame. Meanwhile, the 1080p resolution of the LCD option is a step down.

All of this leaves the latest XPS 13 stuck between an LCD rock and an OLED hard place. We love the OLED screen for its looks, but the battery life hit is awfully hard to stomach. **–MADELINE RICCHIUTO**

BENCHMARKS

	Dell XPS 13 9350 OLED (Intel Core Ultra 7 258V)	Dell XPS 13 9350 LCD (Intel Core Ultra 7 258V)	Dell XPS 13 9345 (Snapdragon X Elite)	Asus Zenbook S 14 (Intel Core Ultra 7 258V)
Geekbench 6 single-core (pts)	2,772	2,660	**2,797**	2,751
Geekbench 6 multicore (pts)	11,033	10,846	**14,635**	11,157
Handbrake time (m:s)	08:28	08:17	**04:41**	08:30
25GB file copy time (seconds)	17.75	**17.5**	20	17.8
Transfer rate (MBps)	1,513	**1,534**	1,342	1,513
Civilization VI: Gathering Storm (1080p) (fps)	**53**	51	48	48
Battery life (web surfing h:m)	08:40	18:34	**19:01**	13:51

Best scores are in bold.

VERDICT

7

Dell XPS 13 9350

➕ **OLE, OLE, OLE, OLED** Fabulous design; Stunning OLED screen; Great performance.

➖ **OH-LED** The OLED screen option absolutely wallops battery life.

$1,899 www.dell.com

SPECIFICATIONS

CPU	Intel Core Ultra 7 258V
GPU	Intel Arc integrated graphics
NPU	Intel AI Boost
RAM	32GB
Storage	1TB
Display	13.4-inch 60Hz 3K (2,880 x 1,800) Tandem OLED
Dimensions	11.62 x 7.84 x 0.6 inches
Weight	2.6 pounds

Asus Zenbook A14

There's something light at the end of the tunnel

AS YOU SWIPE through this magazine, you'll often be confronted by bulky amalgamations of metal and whatever the material they make circuit boards out of is. They're heavy, they have bits that stick out, and they have colored lights and visible cables. They're PCs meant for rendering polygons or housing large language models—not the sort of thing you can slip in a bag and take with you.

The Asus Zenbook A14 is extremely slippable. It's remarkably light at just over two pounds, and even though there's a full-size 14in OLED screen on the front, it feels like it's smaller, like Apple has re-released its 11in line of MacBook Airs. It's that market-sector-creating release that the A14 is targeted at. If there's a Windows machine that can compete with the Air for slimness (the Zenbook is already ahead on weight, as the latest Air weighs 2.7lbs in its 13in form), then it's this.

Inside, you'll find a new, low-end variant of the Snapdragon X processor, with eight ARM cores instead of the 10 on the Plus and 12 on the Elite, and an integrated GPU that runs at about half the power. The X1 may be short on cores, but it's just as feisty as any other Snapdragon we've tested, zipping through CPU-bound work. The Adreno GPU it's attached to is less successful, making it hard to recommend as a laptop for heavy graphics work or gaming (though it might make an excellent streaming device, thanks to its OLED display and full-size HDMI port that you can hook up to a TV without needing an adapter), but that's not really what this machine is all about.

It's meant to be easily portable, and that's something Asus has excelled at with its design. The casing is plastic, and available in gray or a sort of sandy beige. This choice has shaved off the ounces without compromising on its stability, though it does sound a bit hollow if you drum your fingers on the lid, something you wouldn't notice in an aluminum casing. The other thing you won't notice is much noise—as the plastic doesn't do as much for passive cooling as a metal frame would, the A14 uses a couple of lightweight fans and some heat pipes to channel heat away, and there's a vent underneath the machine. It's still incredibly quiet, producing little more than a whisper under normal usage.

However, there's a downside. The Adreno GPU that's shipping with the current generation of Snapdragon chips is weak, though it is compatible with AMD's FSR upscaling, and Qualcomm even has its own solution: Snapdragon Game Super Resolution. This one is particularly feeble. The GPU on this low-end chip can manage about 1.7 TFLOPS, which puts it in PlayStation 4 territory, but it's going to be less well optimized than the console, so we can't expect miracles. In Geekbench tests, the GPU posted an OpenCL score equivalent to a GeForce GTX 750, a chip from 2014 that was bottom of the hierarchy at the time.

Still, even with its lack of graphics power, this is fun to use. It's so small and light, with snappy CPU performance that makes browsing and office apps fly. The battery life—claimed at 32 hours, but more like 20 in reality—is excellent, too, and the A14 goes to show that low-powered machines can be almost as exciting as RTX 5090 cards. **–IAN EVENDEN**

VERDICT

8

Asus Zenbook A14

+ POWER UP Good CPU performance in a thin and very light plastic shell.

– POWER DOWN It's not up to much in terms of graphics processing power.

$1,100, www.asus.com

SPECIFICATIONS

CPU	Snapdragon X X1 26 100
NPU	Qualcomm Hexagon
Graphics	Qualcomm Adreno
Memory	16GB LPDDR5X
Storage	1TB PCIe 4.0 SSD
Screen size	14in
Screen type	OLED
Resolution	1920 x 1200
Refresh rate	60Hz
Ports	1x USB 3.2 Gen 2 Type-A, 2x USB 4.0 Gen 3 Type-C, 1x HDMI 2.1, 1x 3.5mm audio
Wireless connectivity	Wi-Fi 6E, Bluetooth 5.3
Dimensions	12.23" x 8.42" x 0.63"
Weight	2.16 lbs

© UNSPLASH

The Type-C ports are USB 4, and the HDMI is v2.1, which offer lots of connectivity options.

Gigabyte X870 Aorus Elite WiFi7 Ice

Bright, white, and packed with extras

AMD AND ITS MOTHERBOARD partners have recently launched cheaper chipset options for the AM5 socket. But a quick glance at the latest B850 boards sees quite a few retail for more than the Gigabyte X870 Aorus Elite WiFi7 Ice, while also lacking features such as USB4 and high-end audio codecs. Long story short: this is a nicely featured board, given the price tag starts with a '2'.

The X870 chipset is all about making features such as USB4 and PCIe 5.0 for SSDs and graphics cards accessible at more mainstream prices than the X870E. It too is a value play of sorts.

Anyway, apart from the AMD CPU socket and chipset, the Gigabyte X870 Aorus Elite WiFi7 Ice has a nearly identical counterpart in the Intel Z890 camp, making for an interesting comparison. You're definitely getting the better deal with AMD, despite a little less RGB lighting, thanks to a few more ports and power delivery phases. The Z890 option lacks two-pin thermal probe headers included on this board that allow you to hook up coolant probes in water-cooling systems or fine-tune your case's cooling.

You get more Type-C ports on the I/O panel, too, albeit in USB4 guise rather than Thunderbolt 4. There are more PCIe 5.0 M.2 ports as well, but only one of them doesn't steal lanes from your graphics card. It's also great to see features such as power and reset buttons and a POST code display on a board at this price, as they make life easier when testing, benchmarking, or troubleshooting.

As for the unusual HDMI port on the PCB, it's handy for cases like the Y70 Touch that have integrated touchscreens. Normally, you'd have to connect that to a port on the rear panel. With this socket, you can run the cable inside the case, which is neater, although you're limited to 1,920 by 1,080 at 30Hz maximum.

All of Gigabyte's tool-free features are present, with a push-button graphics card release and latches securing the M.2 drives, plus a single-piece Wi-Fi connector. The upper M.2 heatsink, incidentally, is suitably chunky, and cools both sides of your SSD with thermal pads.

Gigabyte's EFI interface, meanwhile, has received a coat of white paint to reflect the white color coding of the board itself. The EFI looks pleasant, and is well laid out, featuring a superb fan control suite while looking a touch less modern than the MSI or ASUS competition, both of which having recently received upgrades.

Gigabyte has done an excellent job in the cooling department, with the VRM's not topping 44 C in our stress test and the chipset not rising above 35 C. The M.2 temperature was low for such an affordable board, with our PCIe 5.0 SSD only reaching 66 C, which was also 6 C cooler than its Z890 sibling achieved while matching the pricier Gigabyte X870E Aorus Pro. Average CPU temperatures in games are surprisingly low, too.

This Gigabyte is slower than other Socket AM5 boards we've tested, though the *Factorio* and *Baldur's Gate 3* results are more competitive. Power consumption is fairly typical; within single digits of other Socket AM5 boards.

So, apart from a snag in its layout, with the top and middle of the PCB almost devoid of fan headers, the Gigabyte X870 Aorus Elite Wi-Fi Ice is decent value. You get good cooling for SSDs and VRMs, plenty of USB ports, and tool-free features to help speed up your build time.

Sure, you might have to make use of long fan cables or extensions if you plan on lots of fans at the top of the case, and be careful how many M.2 drives you fit. Those caveats aside, the Gigabyte X870 Aorus Elite Wi-Fi Ice looks fantastic, offers a decent, balanced set of features and should keep on delivering long after the AM5 socket itself is sunsetted, probably around 2027. **–ANTONY LEATHER**

VERDICT

8

Gigabyte X870 Aorus Elite WiFi7 Ice

+ ICE, ICE BABY Great cooling; Decent value; Good feature set.

− PUT IT ON ICE Imbalanced fan header placement; Two of three M.2 slots steal lanes.

$290, www.gigabyte.com

SPECIFICATIONS

Socket	AMD Socket AM5
Chipset	AMD X870
CPU compatibility	AMD Ryzen 7000/8000/9000 desktop
Form factor	ATX
Memory support	DDR5-4800 to DDR5-8200 (OC), up to 256 GB
Storage	4x M.2, 4x SATA
USB (rear)	2x USB 4 Type-C 40 Gbps, 2x USB 3.1 Type-A 10 Gbps, 4x USB 3.0 Type-A 5 Gbps, 4 x USB 2.0
Display	1x HDMI 2.1, 2x USB/ DisplayPort
Networking	Realtek 2.5G LAN, Wi-Fi 7
Audio	Realtek ALC1220

BENCHMARKS

	Gigabyte X870 Aorus Elite WiFi7 Ice	Gigabyte X870E Aorus Pro	MSI MAG X870 Tomahawk Wi-Fi
Cyberpunk 2077 (fps)	111	111	**117**
Baldur's Gate 3 (fps)	100	96	**103**
Gaming power draw (W)	**109**	112	114
Cinebench R24 multi-thread (pts)	**1,836**	1,786	1,820
Blender 4.2.0 (samples per minute)	**150**	147	**150**
Peak VRM temperature (C)	**44**	46	49
Peak SSD temp (C)	**66**	**66**	72

Best scores are in bold. * All boards tested with an AMD Ryzen 9 9900X CPU and 32 GB Corsair Vengeance Pro DDR5-6000

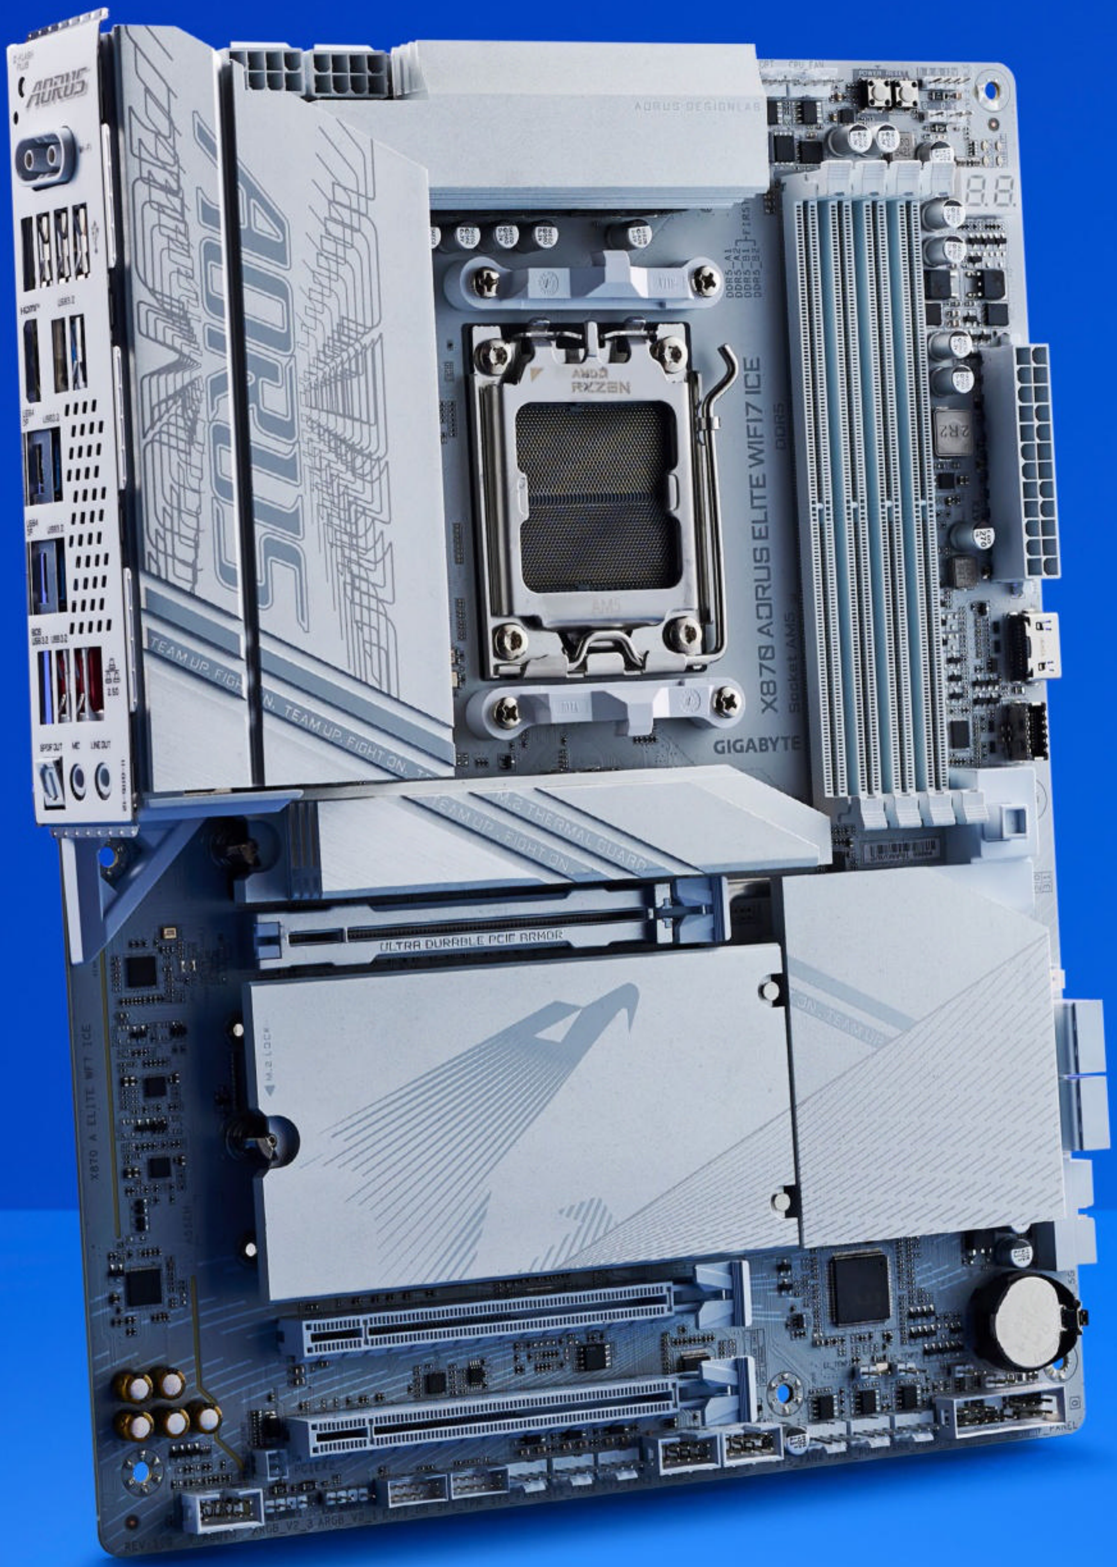
X870 AORUS ELITE WIFI7 ICE
DDR5
Socket AM5
GIGABYTE
AORUS
TEAM UP. FIGHT ON.
M.2 THERMAL GUARD
ULTRA DURABLE PCIE ARMOR
AORUS DESIGNLAB

There are three PCIe 5.0 M.2 ports, but only one won't steal lanes from your graphics card.

Logitech G RS Wheel Hub and Track Wheel

The wheel upgrade we've been waiting for

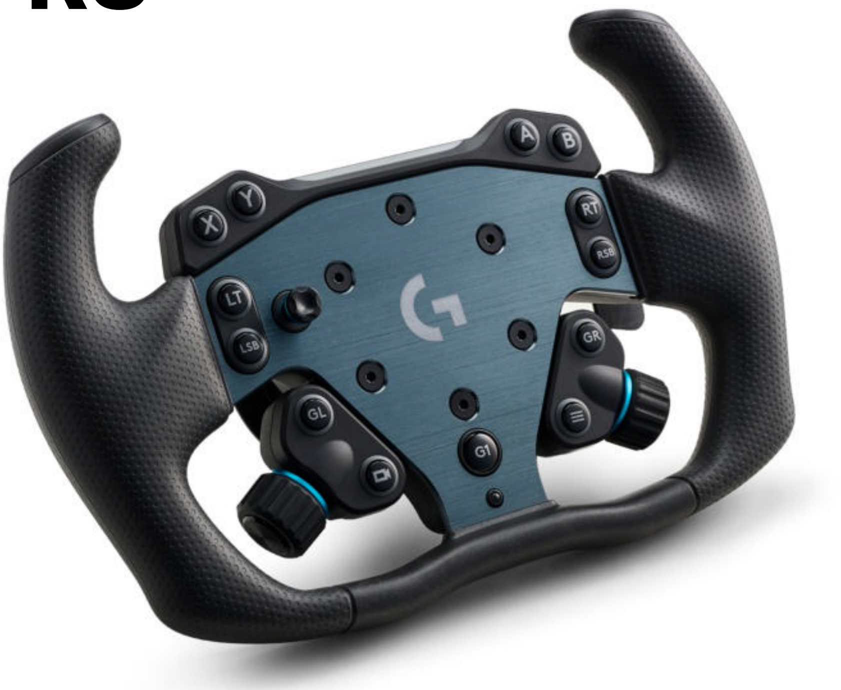

OWNERS OF THE Logitech G Pro Racing Wheel, rejoice—we can finally upgrade our sim racing setup as we always hoped, with new shifters, handbrakes, wheels, and wheel hubs. Logitech has two new racing wheels on offer: one designed for the track, and another for the rally—or even classic car—crowd.

There are two parts to this latest release: the RS Wheel Hub containing the interface, paddles, and electronics to attach the wheel to the Pro Racing Base, and either the RS Track Wheel, as reviewed here, or RS Round Wheel, which are simply metal and a rubber-like plastic. You can use any pedals, shifter, or handbrake with the Logitech gear, provided you plug it into your PC separately. The Logitech stuff all combines into a single USB cable via the USB hub on the rear of the Pro Racing Base/Wheel.

Anyway, the RS Hub contains 15 buttons, one-directional stick, two dials, and two rear paddles. Compared to the Pro Racing Wheel, the RS Hub offers two extra buttons and a better overall layout with symmetrical dials, which are easier to use while racing. However, it disappointingly lacks the dedicated clutch paddles.

While only useful for launches at the very start of a race, F1 cars use a clutch paddle rather than a pedal. As an F1 fan, I was hoping to use the RS Track Wheel for a more authentic experience. The RS Track Wheel does make it up in other ways, at least; rounding off the RS Wheel Hub is an RGB LED strip that corresponds to your revs in-game. This actually takes over from the rev meter already included on the Pro Racing Base as standard.

Onto the RS Track Wheel. It's a hunk of metal over 3mm thick and wrapped in a molded TPE grip. It feels like it'll last a good few thousand laps, at the very least, but it does lack the softer, more premium feel of the Pro Racing Wheel, which comes wrapped and stitched.

Combining the RS Wheel Hub and RS Track Wheel couldn't be easier. Slot the wheel over the hub, and it fits snugly into place. Then go around attaching the six included bolts with the larger of the two included Allen keys. Once combined, the RS Wheel Hub and RS Track Wheel comes in at 1,220 grams versus the 1,451 grams of the standard Pro Racing Wheel.

BURNT RUBBER

So, the Track Wheel weighs a good bit less, and combined with that smaller size and open-wheel style design, you can really throw it around at high speed. It's agile, but with the option to fully grip either side to steady it during a tough corner, with 11nm of torque pushing through it. It's also wicked fun.

Long story short, I've had a blast with the Track Wheel installed. It just feels right to be racing around in F1 24 with a compact wheel like this. Easing into the deep corners with a full rotation, like the pros, is really satisfying. I've also taken it for a spin in *Forza Motorsport* and something altogether more absurd: an old Lotus F1 car at Brands Hatch. That thing goes. The grip on that thing is otherworldly, and I could nail the speed through the corners with a little tweak of the Track Wheel either way.

Altogether, I'm itching to get back into the racing seat even as I write this. Considering this is a competitively priced unit compared to even Thrustmaster's cheaper F1-style wheels, I'm thoroughly impressed with it. Pulling it out of the box, it feels like all the premium product you want it to be, and maybe a little more for the money. The lack of clutch paddles is a bit of a shame, but this is undeniably a great upgrade for a sim racing setup.
–JACOB RIDLEY

VERDICT
9

Logitech G RS Wheel Hub and Track Wheel

+ **WHEELY GOOD** Premium feel; Competitively priced; Feels great in-game.

− **TURN OFF** No clutch paddles; Doesn't work with Pro Racing Wheel hub.

$200 (combined) www.logitech.com

SPECIFICATIONS

Buttons	15
Connectivity	Quick release attachment
Features	2x paddle shifters, RGB rev lights, 2x dials, 1x directional stick

NZXT Lift Elite Wireless

Light and lovely, there's just something missing to tip it into top-tier territory

IF THERE'S ONE BIG TREND in the world of PC peripherals that most folks can get behind, it has to be the ultralight gaming mouse. We've seen options emerge from the likes of Logitech and Glorious as some of the best gaming mice out there. The latest to give it a go is NZXT with its Lift Elite Wireless.

As with the other peripherals released in this new Elite range, the Lift Elite Wireless is targeting an aggressive price point at $89. With the spec sheet on offer including a 26,000 DPI sensor, 57g weight and support for up to 4K wireless and 8K wired polling rate, that makes this quite the compelling option against the Logitech Pro X Superlight 2 Dex, and especially the Razer DeathAdder V3 Hyperspeed, our top contender.

Unlike the more ergonomic and contoured lines of the DeathAdder, the Lift Elite Wireless features a more symmetrical shape to its lightweight chassis that's reasonably comfortable to hold. It fits my medium-sized hand well enough, although folks with larger hands may want to seek a bigger option. The plastics used feel solid, and come with a pleasant texturing to aid comfort.

Despite the 57g weight, the Lift Elite Wireless doesn't feel cheap, as some ultralight options can. Button placement is sensible for the most part, with the main options and a scroll wheel on top and two navigation buttons on the left. The button to change DPI is on the underside, and owing to the open underside, there isn't anywhere to store the USB-A receiver needed for connectivity.

That light frame makes the Lift Elite Wireless quite the zippy mouse to use in competitive titles such as *Counter-Strike 2* and *Apex Legends*. This NZXT option also benefits from a 26,000 DPI sensor for especially fast movements. While the competition can go even higher, it doesn't make much of a difference to mere mortals like you and I.

That high DPI nonetheless made sweeping movements in *Counter-Strike 2* an absolute breeze for when I needed to swing around to fire at enemies in a flash, although it is too jittery for productivity workloads. By default, it's set to a much more reasonable 3,200 DPI as the maximum, so you will need to bump it up to 26,000 DPI in software, which is easy enough.

The big kicker alongside that high DPI is that the Lift Elite Wireless supports up to an 8,000Hz polling rate. This makes it part of an exclusive club of gaming mice that go above 1,000Hz. This potentially enables lower latency. A higher polling rate means a device reports its position and inputs more frequently (hence Hz as the measurement). 4,000Hz is possible over its bundled 2.4GHz receiver. The connectivity here is pretty simple, too, via either the 2.4GHz receiver or USB cable. There isn't support for multipoint Bluetooth for added versatility.

The 70 hours of rated endurance is reasonable, although sits well behind the DeathAdder V3 Hyperspeed's 100 hours. It's worth noting that figure is with a 1,000Hz polling rate. With anything higher, expect the battery life to drop. I saw around 20-30 hours of playtime at the maximum 4K wireless polling rate.

Against the competition, the Lift Elite Wireless' software customization feels a little basic. It's all present in one tab, which is convenient, although you only get options for configuring the five DPI levels, polling rate, lift-off distance, and button remapping. You also get the option to save up to five profiles. There aren't any advanced surface calibration or game-specific profiles, as is found in Logitech's G Hub for instance.

This is a good mouse, but it doesn't match the Razer's specs and performance, or the Logitech Pro X Superlight 2 Dex's clever software. In short, this doesn't do anything wrong, but not even the stylish white colorway is enough to give it that competitive je ne sais quoi. **–REECE BITHREY**

VERDICT

7

NZXT Lift Elite Wireless

⊞ **FANCY A LIFT?** Competitive pricing; Good performance; Decent battery life.

⊟ **TAKEN FOR A RIDE** Slightly basic software; Doesn't do anything to really stand out.

$80 www.nzxt.com

SPECIFICATIONS

Buttons	Six
Feet	PTFE
Connectivity	2.4 GHz Wired (USB-A to USB-C cable)
Sensor	PixArt PMW3395
Max DPI	26,000
Max acceleration	50 g
Max speed	650 IPS
Polling rate	4,000 Hz (wireless)/ 8,000 Hz (wired)
Battery life	70 Hrs
RGB lighting	None
Warranty	Two years

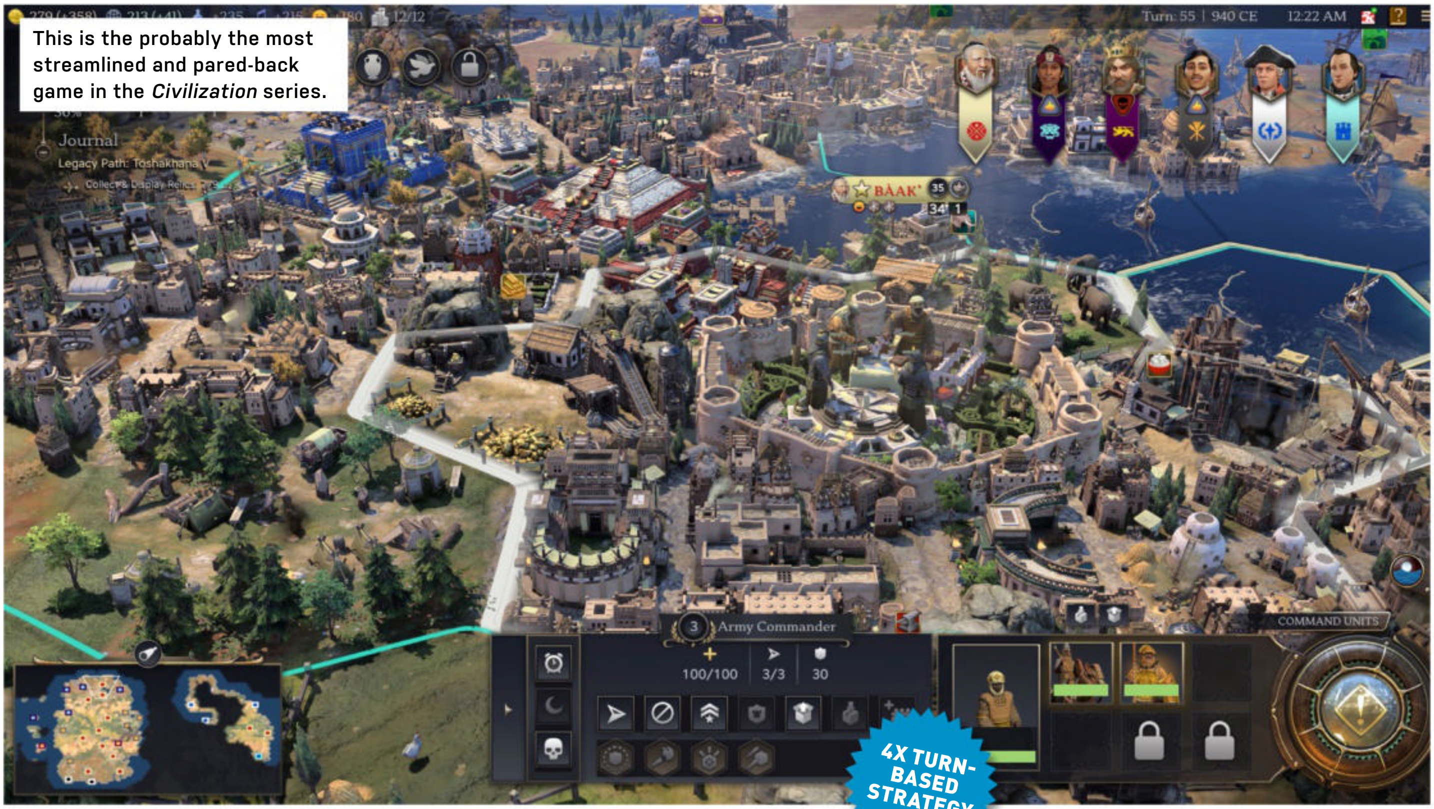

This is the probably the most streamlined and pared-back game in the *Civilization* series.

4X TURN-BASED STRATEGY

Civilization 7

Lighter, leaner, and with sights set on a new audience

MORE LIKE a series offshoot than a *Civilization 6* sequel, *Civilization 7* takes the venerated grand strategy series in a different direction from the past two games. Certainly, it's the most streamlined and pared-back the series has been in a long time. But as we'll see, old *Civ* hands might yearn for the depth and diversity of its predecessors.

Founded in the pre-SVGA age of gaming antiquity, *Civilization*'s randomized yet celebratory take on human history has stood the test of time. Over its first several entries, the series rapidly evolved to better encapsulate the vibrancy and madness of human history in a turn-based strategy format. But in more recent outings, it has become more about rejigging rather than revolutionising that cavemen-to-cosmos saga.

The boldest idea that *Civ 7* introduces are Age transitions. Leaders and civilizations have now been separated. You keep the same leader throughout a game, but your civ will change. Each game is split into three Ages—Antiquity, Exploration, Modern—and each time you move into a new one, time skips several hundred years, and you pick from one of a bunch of new civs to lead, specific to that era.

That age transition is, by design, far from seamless. As you near the end of an Age, which all civs collectively reach by making progress down four linear Legacy Paths, the world will face some kind of crisis, such as plague or barbarian invasions. Once the crisis reaches its head, the current iteration of your civilization collapses, and you pick a new one to build over its foundations.

There's a hint of the roguelike about it all, as your foot-up in the next era is built on your efforts of the previous one. From a pacing perspective, this works wonders. At the point in previous games where things would start getting a bit stodgy in terms of unit clutter, city micromanagement, and losing players basically being out of the victory race by the midway point, *Civ 7* speeds things up and narrows the gap between the haves and have-nots.

Meanwhile, *Civ*'s age-old system of citizen management, whereby you could manually assign a city's population to work specific tiles as a way of adjusting a city's focus, has been replaced by slottable resources. In my relatively short time with the game so far, I've even managed to find some dirty little tricks with resource-slotting, like adding several camels to a city to increase its resource capacity, then building the Tomb of Askia Wonder to grant it extra gold and production for each resource in that city.

DEVOLUTION

While the idea of distributing resources throughout your empire to where they're most needed is a sensible one, by the late game, I had gotten a bit fatigued with the resource screen, which winds up containing dozens of little resource icons to manually micromanage. What's more, a little too often, *Civ 7*'s fat-trimming cuts into the lean part of the meat as well, losing some richness and flavor in the process. Diplomacy, for instance, feels very thin, and interactions with other leaders a little too transactional.

This iteration of *Civilization* is also weirdly, well, civil. City-states and barbarians from previous games have been rolled into Independent Powers—some start off hostile and effectively act as barbarians, but all can be befriended

© 2K GAMES

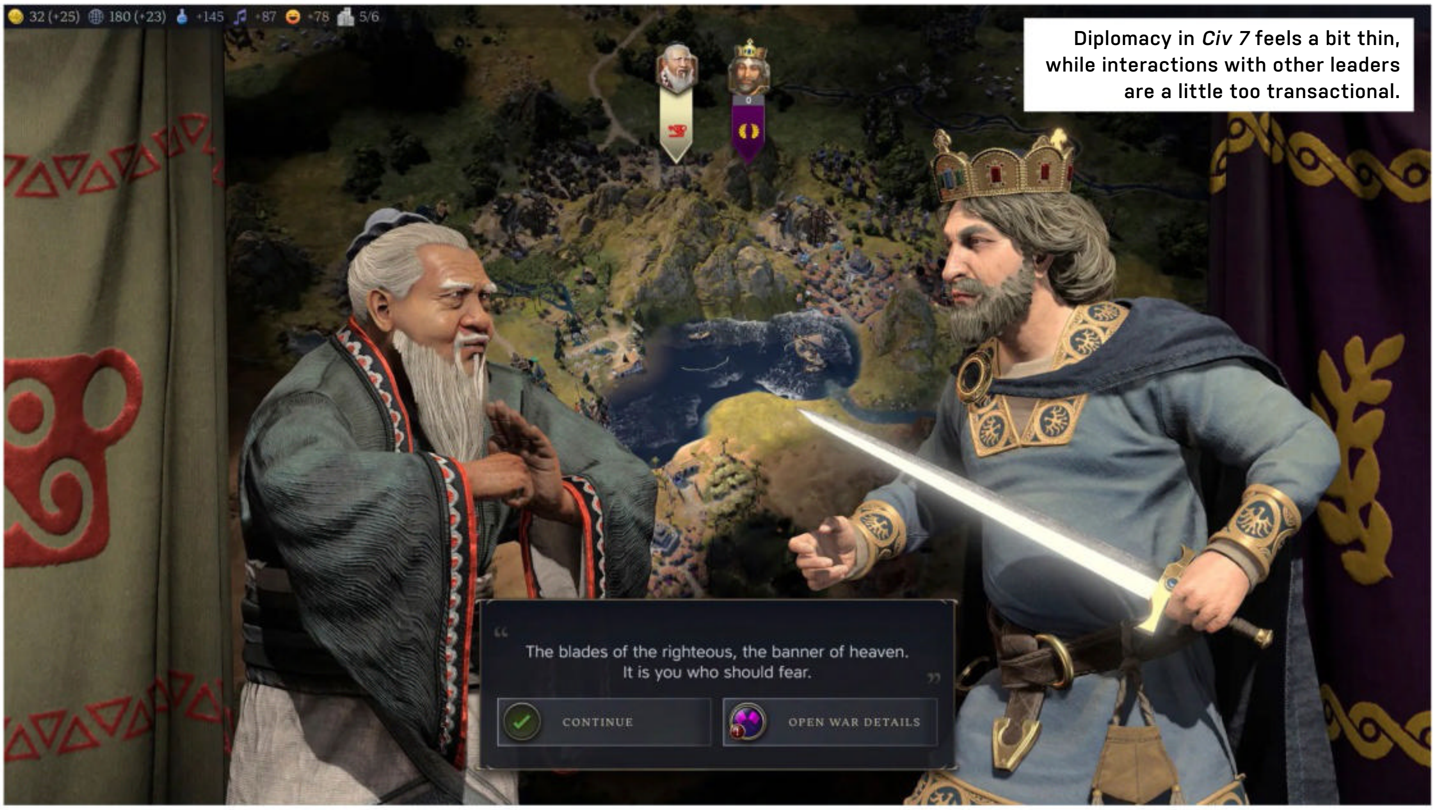
32 (+25) 180 (+23) +145 +87 +78 5/6
CONTINUE
OPEN WAR DETAILS

"The blades of the righteous, the banner of heaven. It is you who should fear."

Diplomacy in *Civ 7* feels a bit thin, while interactions with other leaders are a little too transactional.

by spending enough Influence (a new currency used for every diplomatic interaction). Once you start the befriending process, you're guaranteed to become their suzerain for the rest of the age, unless they get wiped out.

To be clear, most of *Civ 7*'s shortcomings are relative to the two very fleshed-out games preceding it. Indeed, *Civilization 7* feels like a reaction to the maximalism of its predecessors—sleeker and speedier, colder and less complex. But while it has undoubtedly addressed some longstanding issues, there's too much that's been stripped away to call this meaningful progress. **–ROBERT ZAK**

VERDICT 7

Civilization 7

CIVILIZED More accessible than before; Age transitions are great for pacing.

LOBOTOMIZED Interactions and diplomacy dumbed down; A touch too much has been stripped away.

RECOMMENDED SPECS CPU, Intel Core i5-10400 / AMD Ryzen 5 3600X. GPU, NVIDIA RTX 2060 / AMD RX 6600 / Intel Arc A750. RAM, 16GB. OS, Windows 10.

$69.99, civilization.2k.com, M-rated

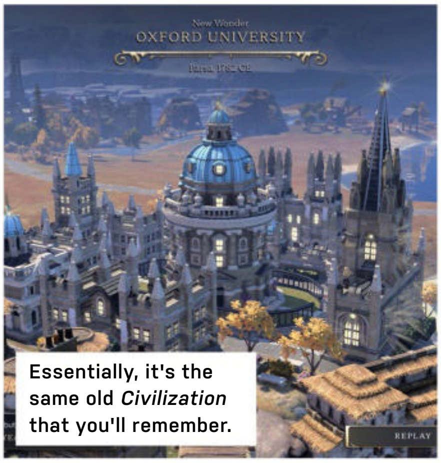
New Wonder
OXFORD UNIVERSITY
REPLAY

Essentially, it's the same old *Civilization* that you'll remember.

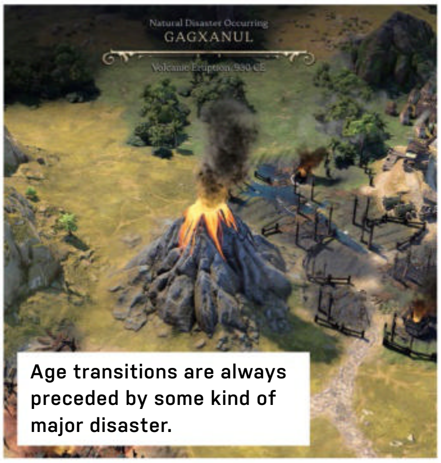
Natural Disaster Occurring
GAGXANUL
Volcanic Eruption, 930 CE

Age transitions are always preceded by some kind of major disaster.

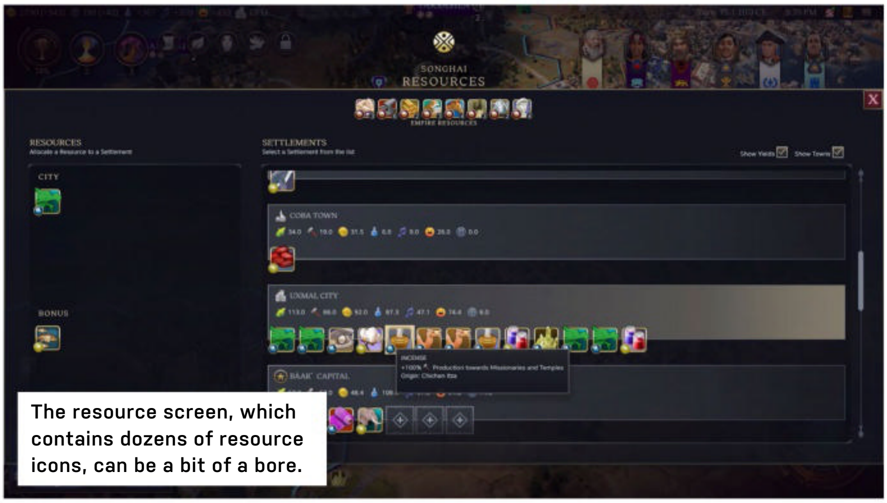
SONGHAI
RESOURCES
EMPIRE RESOURCES
RESOURCES
SETTLEMENTS
CITY
BONUS
COBA TOWN
UXMAL CITY
BAAK CAPITAL

The resource screen, which contains dozens of resource icons, can be a bit of a bore.

HEAD TO HEAD
WINNER
★★★★★

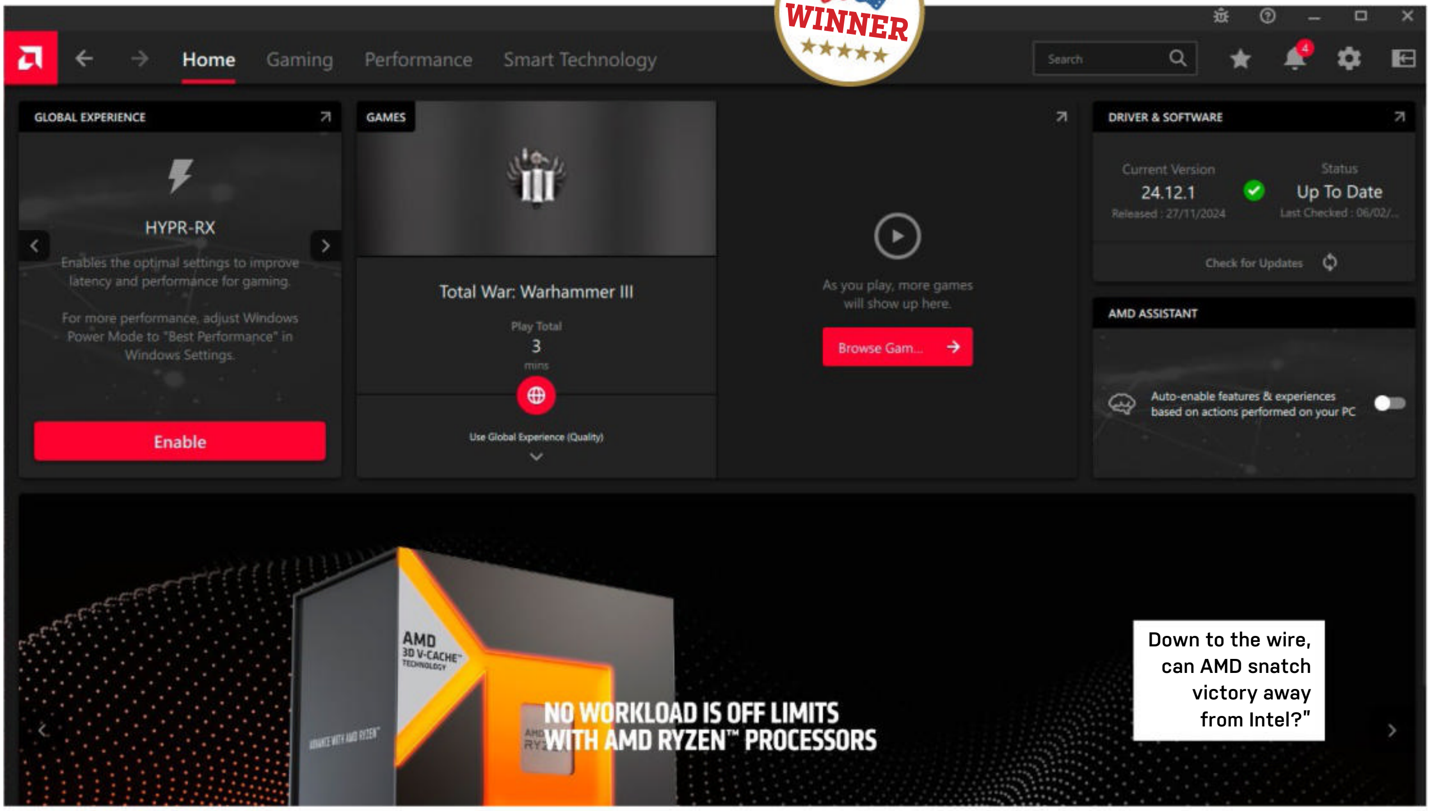
Home
Gaming
Performance
Smart Technology
Search
GLOBAL EXPERIENCE
HYPR-RX
Enable
GAMES
Total War: Warhammer III
Play Total
3
mins
Use Global Experience (Quality)
Browse Gam...
DRIVER & SOFTWARE
Current Version
24.12.1
Released : 27/11/2024
Status
Up To Date
Last Checked : 06/02/...
Check for Updates
AMD ASSISTANT
NO WORKLOAD IS OFF LIMITS
WITH AMD RYZEN™ PROCESSORS
AMD
3D V-CACHE™
TECHNOLOGY

Enables the optimal settings to improve latency and performance for gaming.

For more performance, adjust Windows Power Mode to "Best Performance" in Windows Settings.

As you play, more games will show up here.

Auto-enable features & experiences based on actions performed on your PC

Down to the wire, can AMD snatch victory away from Intel?"

© AMD

AMD Software: Adrenalin Edition vs Intel Graphics Software

Let the battle continue

LAST ISSUE, we promised to test graphics software from the big three, and that's what we're delivering. Last time, we put both Nvidia's freshly released 'NVIDIA App' (still hate that capitalization) and AMD's Software: Adrenalin Edition into the *Maximum PC* labs to figure out who exactly was the king of the hill when it came to graphics setting tweaks.

In perhaps a surprise to very few, AMD came out on top, slaughtering the green competition with an advanced suite of metric logging, performance tweaking, and better laid-out software than Nvidia could match—at least not yet. That's not a surprise, because the Nvidia App relaunch is relatively new, and still under development. AMD Adrenalin has been iterated on time and time again since the Radeon 5000 series, and has developed some serious chops.

It needed it, too, with little in the way of competitive top-end graphics cards and Intel nibbling at its heels on the low end, having an impressive graphical suite was tantamount to giving it a slight edge in the field. Having slayed Nvidia, can AMD do the same to Intel?

Team Blue is relatively new to the scene. We're effectively on its second generation of cards, with the Intel Arc B580. As a series, Intel's cards represent dominance when it comes to a price-to-performance metric. Simply put, it stripped the crown from AMD's shiny red head, hoiked it into a chest, and buried it on a desert island somewhere; its GPUs are that impressive.

Of course, with a line of fresh graphics cards out of TSMC's foundries, it needed a solid software solution to back it up. In many ways, similar to Nvidia's App, it's still under development, and you can tell that from the outset. Still, it's a seriously slick-looking piece of kit, and one that makes jumping ship to Intel less painful, particularly if you care about making sure your hard-earned cash is invested well.

SAPPHIRE SUCCESS?

To start with, Intel's Graphics Software has an absolutely stunning, beautiful, crisp, and clean UI. You get those signature blue colors, of course, with a clean, legible design, font, and all the info you need exactly where you need it. The home page clearly displays the current driver, whether you're up to date or not, and when the last check was performed.

More than that, there's also a number of driver download options here; for one, you can immediately control how you get software updates, from just plain notifications, to complete download and install (or disabled entirely). If it's in its auto 'download and install' mode, you can even tell Intel how you want the installation to go, either via your traditional iterative 'Express' update (where it just installs over the top retaining settings) or a 'Clean' mode (where it'll remove the old driver and settings and install the new driver fresh).

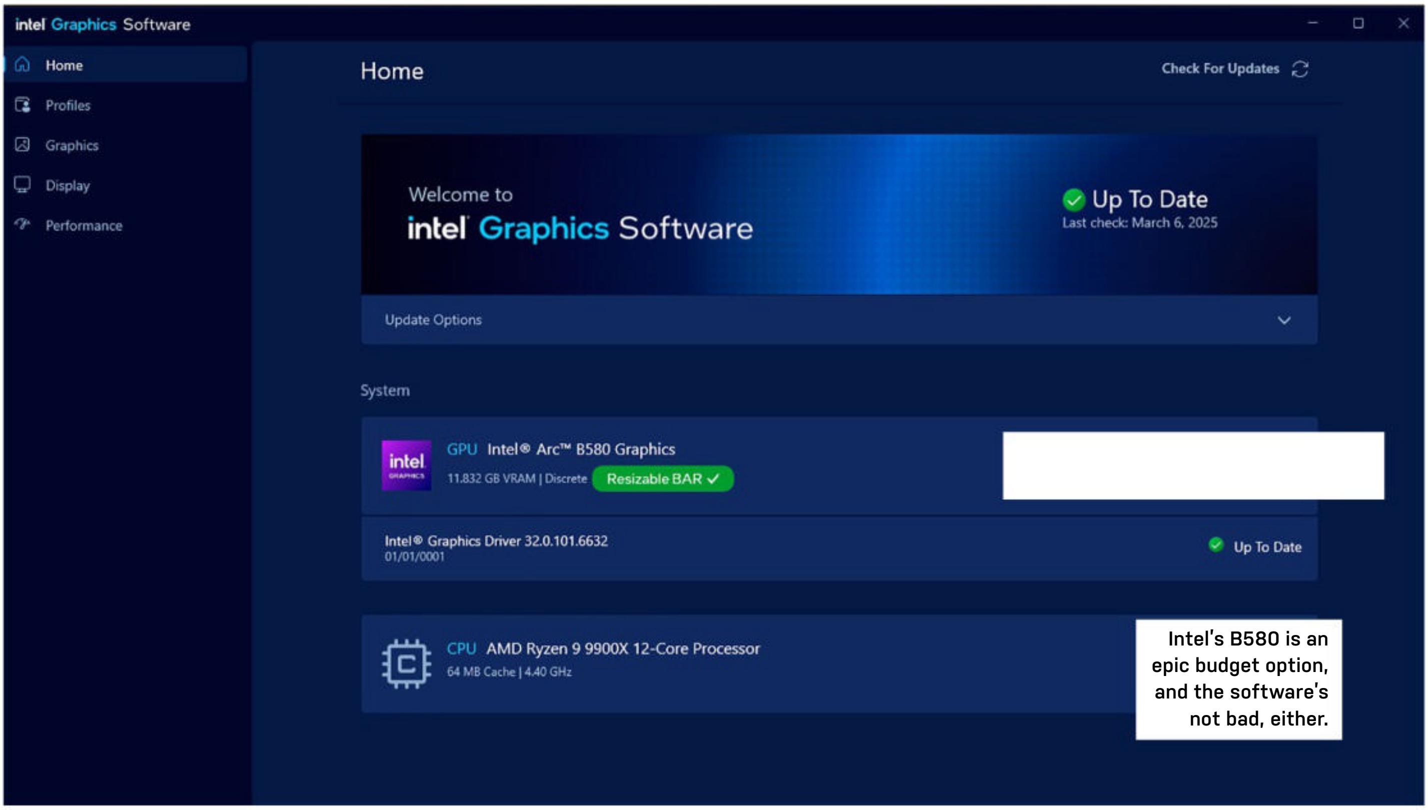

✓ Up To Date
Last check: March 6, 2025

GPU Intel® Arc™ B580 Graphics
11.832 GB VRAM | Discrete

Intel® Graphics Driver 32.0.101.6632
01/01/0001
✓ Up To Date

CPU AMD Ryzen 9 9900X 12-Core Processor
64 MB Cache | 4.40 GHz

Intel's B580 is an epic budget option, and the software's not bad, either.

© INTEL

More impressively than that, you can check a box to ensure you only get WHQL-certified drivers. That's Microsoft's Windows Hardware Certification Program, ensuring drivers pass its strict software requirements, which is great if you prefer stability over the latest performance boost.

Moving on to Profiles and Graphics, these are fairly consistent with what you'll find with AMD and Nvidia. In profiles, you can configure settings per app. For instance, you can set up low latency mode, frame sync choice, or adaptive tesselation settings per program. It's limited compared to AMD's offerings, but similar to Nvidia in its overall breadth, at least in the app (Nvidia's Control Center is more advanced, but looks like it's stuck in Windows 98, and is technically a separate application). The 'Graphics' tab takes those same settings and allows you to apply them globally.

Under the Display tab, it allows you to control each screen that the GPU can detect. There's no resolution or frame rate control (there is a link to the Windows settings, but that's it). Nvidia still has the edge, as you can position the screens in a virtual setup directly in the Nvidia App (something AMD also lacks).

What is much better is the live control on hue, brightness, contrast, and more. All of it can be directly tweaked per panel and changed on the fly. Intel has also included VRR support options, allowing you different levels of control, plus color-depth control, shifting from 8-bit to 10-bit, and color format too, shifting from RGB to YCbCr 4:2:2 or 4:4:4. All of that is configurable per monitor.

The Performance tab is where you'll find system monitoring and a Tuning tab. This is a bit of a mix of the Green and Red's (brown?) solutions here. Stats are broad, and there are some awesome frame-rate monitor metrics, but no way of logging this. Nvidia does have Frameview as a separate entity, but with AMD, you can log directly. Where Intel has the edge of over Nvidia is in the tuning department, where you'll find control over your GPU, from voltage to power limit to even the frequency curve over milivolts. VRAM control is there, too, although it's heavily simplified, with adjustments only possible for 'total bandwidth' rather than frequency and voltage. There are no profiles—just a 'save and your done' style of overclocking.

Intel sits somewhere in the middle with its Graphics Software. It's clean, slick, and gets you all the relevant information and driver updates, and some neat features. It lacks the detail of AMD, yet edges out Nvidia on the software front. Of course, this is all a fairly moot point; it's not the graphics software that really holds much sway in what your next GPU is going to be. Still, while Nvidia flounders slightly and AMD's 9070 line rocks the boat, now might not be a bad time to shift allegiances. The days of poor graphics software, middling driver support, and bloated GPU settings seem to be coming to an end. You won't be disappointed if you do decide to go the AMD route.

And with that, our two-parter is complete: Intel, AMD, and Nvidia, all wrapped up. **–ZAK STOREY**

VERDICT
9

AMD Software: Adrenalin Edition

⊞ **ATI RADEON 9700 PRO** Outstanding feature set; Multiple tuning options; Undervolting; Per-program settings and options; Logging features built in; Most features built in directly.

⊟ **AMD RADEON RX VEGA 64** Only available with Radeon; Lacks display control in app; Occasionally difficult to install.

Free, www.amd.com

VERDICT
8

Intel Graphics Software

⊞ **iCUE WOULD BE JEALOUS** Clean UI; Easy to use; WHQL-drivers; Display tuning is great; Overclocking is strong; Competes on almost all featuresets with the big two.

⊟ **IT'S PRONOUNCED JIF** Lacks overclocking profiles; No logging available; Display features could be better.

Free, www.intel.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Retro rewind
> 8K reveries
> Virtual realities

New laptop, old-school gaming

Hello *Maximum PC*. I'm a long time reader of the mag in its many forms over the years, and as that may suggest, I'm also a bit on the older side. I used to be a big gamer—we're talking DOS and the 16-bit era—but I've fallen out of love with keeping up with hardware demands, and all the online 'live service' stuff just doesn't appeal to me.

But I still like to play the odd game from time to time, and even recently bought an Xbox-style controller to pair with my Dell XPS work laptop. Part of the inspiration was a friend who has told me about game emulation, and the ability to play Super Nintendo and Amiga 500 games on PC. I'd never heard of this before and it sounds great! How does it work? Can I get this sort of thing going on my simple work laptop? I'd love to fire up *Monkey Island* again and relive some of those great gaming memories.

All the best and keep up the good work. **–John Plop**

EDITOR, GUY COCKER, RESPONDS: John, you are about to go down a very, very fun rabbit hole, and I can't wait to help you.

Emulators, as the name suggests, use your modern PC hardware to recreate gaming experiences of yesteryear—often thanks to the hard work of individuals who have reverse-engineered game, BIOS, and OS code from old systems.

The results can be astonishing—not only do you get to play stuff from gaming's golden ages; they're often improved experiences, too. You can use modern hardware to improve the resolution, frame rate, and performance of many older titles, while also introducing modern conveniences like save states for games that didn't have them upon release.

Legally, it's a bit of a gray area. Emulators themselves can usually be built, distributed, and used without issue—though some game companies, notably Nintendo, have taken umbrage at emulators built around current systems, and sent their legal eagles at their developers. Yes, you can emulate even relatively new consoles like the Nintendo Switch. The more difficult area is acquiring 'ROMs'—the code dumps of the games themselves. Distributing them is not legal, and though there are methods of pulling these off of original cartridges and disks, more often they're shared online. Proceed with caution.

Once you're over those hurdles, the sky is the limit, and even a very basic, entry-level laptop can emulate games of the era you refer to. Start with the Retroarch emulator—it uses 'cores' in one program that let you play multiple consoles from inside one interface. There's a learning curve, but its online documentation is great, and can be tailored to your needs.

For your specific *Secret of Monkey Island* request, for instance, you could use Retroarch to get the adventure game going—or it might be simpler to use the ScummVM standalone emulator. Both are free, and Retroarch is using ScummVM in its multiplatform UI.

If setting up these emulators sounds too challenging, look at LaunchBox (PC) or Batocera (Linux). These 'front ends' take a lot of the busy work out of configuring emulators, and have jukebox-like interfaces that show game artwork, video previews, and even game manuals when fully set up.

Don't be turned off of modern gaming just because of that PC hardware arms race. Now you've got an Xbox controller, head to Valve's Steam PC gaming store, and check out some of the lightweight indie titles. Some of the best games of recent years—*Balatro*, *Animal Well*, *Into The Breach*—have very low system requirements.

If you do want to explore some of the more graphically intense PC games of more modern times, don't rule them out entirely. Providing you've got a half-decent broadband connection, there are ways to play them on even modest devices.

Services like Nvidia GeForce Now and the premium tier of Xbox

↘ submit your questions to: editor@maximumpc.com

Game Pass let you stream triple-A games over the internet—a bit like Netflix, except you control the action with your gamepad. These services have gotten so good in recent years that the untrained eye wouldn't be able to tell the difference between them and local, native gameplay. You pay for a subscription on a monthly basis, so you can try if it's for you for around ten dollars without committing long term. Definitely worth a go!

Happy gaming!

5090 black screen blues

With a little help from a buddy and a bucketload of cash, I've managed to get my hands on the MSI GeForce 5090 32G Gaming Trio card. I plugged it in, booted it up, and my love for it was almost instant as it just raced through even demanding ray-tracing stuff. Jumping from a 3070Ti, it's blown my mind.

But I've got a bit of a panic brewing in that in the last few days it's just black screening now almost at random, and sometimes won't come back unless I do a full system restart. I can't return this thing as it came from a brick and mortar US store, and I live outside the US. Have I made an error the size of a month's pay packet?

–L. Hiller,

EDITOR, GUY COCKER, RESPONDS: Well done for managing to nab one of these highly sought-after cards, and even bigger congratulations to your stateside pal who managed to a) find one of these in person and b) get it shipped safely abroad!

you're not alone—we've seen a few reports of users suffering from black screen issues with the new top-tier Nvidia cards, and we imagine it'd bring on a sweat if you've coughed up all that cash.

© LAUNCHBOX

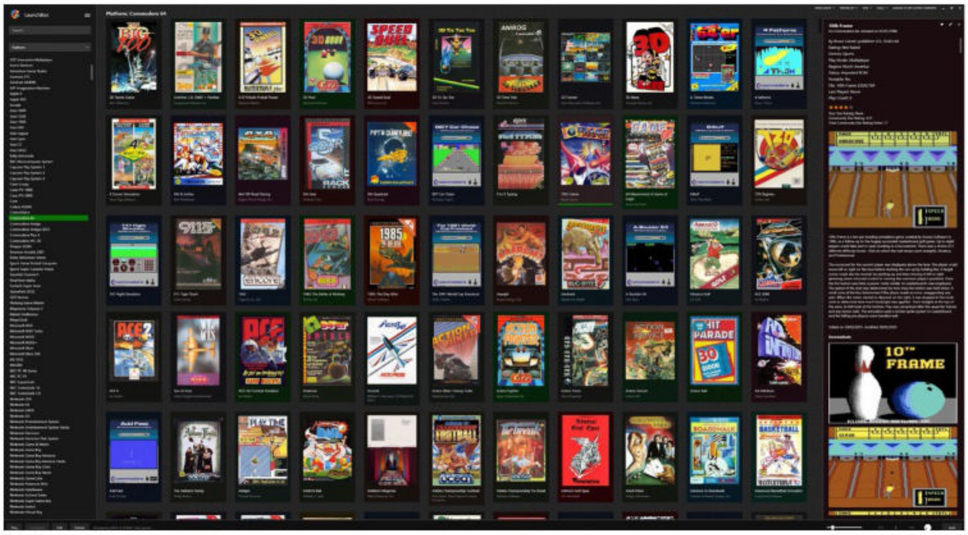

LaunchBox is like a jukebox for retro game collections.

But fear not! It all seems to be down to a VBIOS issue, and a new driver update rolling out should sort this problem out. Just head over to the MSI Center software suite, install the latest driver, and you should be good to go.

All the best.

8K now or never?

For years I've been putting off getting an 8K TV or monitor. I'm a sucker for sharpness, and though 4K has served me well for many years, I get tempted by the thought of even more pixels. With 2025's new GPUs and CPUs looking pretty powerful is now the time to finally pull the trigger?

–Calvin Morrison

EDITOR, GUY COCKER, RESPONDS: Calvin, if this were a two-way conversation, I'd be asking a few more questions to really answer your query. What are you planning to do with your supersharp monitor? Are you in the creative industries, like a videographer or designer? Are you looking to play lots of games? Or are you simply looking to do some office suite work and browse the web?

8K has its benefits, namely screen real estate at higher scaling levels, but performance is dependent on the programs you're using and the rig you're pairing the monitor with. Photo and video editing apps are well optimized for 8K workflows, and the benefits of having that extra clarity when working on visual projects is clear. But top-tier games struggle, even with new visual enhancing 'cheats' like DLSS and frame generation. Simple office applications won't be much of a problem at 8K resolutions, but would I lay out that expense for a high-res spreadsheet? In 9 out of 10 scenarios, I'd say save your cash.

Virtual indecision

What's your take on virtual reality headsets? Meta's models get great reviews, but how are they for PC VR? Would I be better off getting something like the Valve Index?

Thanks for your advice!

–Pádraig Folan

EDITOR, GUY COCKER, RESPONDS: If it were my money today Pádraig, it'd be going on a Meta Quest 3 headset. Though lots of the marketing talks about it being a standalone headset with built-in apps and games, it's also a capable PC VR device. Whether you're hooking up to a PC over a USB-C connection, or streaming over the wireless Air Link application, it's easy to get the Meta Quest working with a PC. You're not just tied to the Meta apps store, either—it's easy to get it working with third-party apps, as well as platforms like Steam.

The Meta Quest 3 gives you the flexibility of wire-free VR when not by your PC, and doesn't sacrifice visual fidelity, with 2,064 x 2,208 pixels per eye. The fact that it's a fraction of the price of the Valve Index, and doesn't need half as many wires or base station motion tracking sensors, seals the deal. You can slash the cost further if you opt for the Meta Quest 3S—still amazing, the concession being a lower-resolution display.

One potential spanner in the works: Valve is rumored to be making a standalone headset, codenamed Deckard. If it's as good as what it produced with the Index when that headset launched in 2019, Meta will have some stiff competition. It might be worth waiting a few months to see if that particular rumor bears fruit. ⏻

THE BUILDS

THIS MONTH'S STREET PRICES...

ALL QUIET ON THE WESTERN FRONT. The calm before the storm. Take your pick, but right now, as we write this, the US government's tariffs have only just come into effect. That's big news for Blueprints and PC building in general, and particularly for us moving forward.

Chinese goods have just been hit with a 20 percent tariff on all exports. That'll extend to the majority of case manufacturing, along with peripherals, monitors, and tertiary PC componentry. There's also been mention of tariffs of up to 100 percent hitting Taiwan, predominantly aimed at crippling TSMC's international supply chain, although those aren't live at this moment.

Given that Qualcomm, Apple, AMD, Nvidia, and Intel all utilize TSMC for their processors, we can expect price hikes on any new stock landing in the US. That said, TSMC does have some capacity to manufacture those chips abroad now, with the Taiwanese government recently announcing the company is now free to manufacture its 2nm products outside the country. Still, that will be a severely limited supply chain heading into the US.

China, Mexico, and Canada have all announced retaliatory tariffs, and the current US government has promised escalatory tariffs if those are enacted. All this is bad news for us custom PC folk, particularly given we're already in the midst of a GPU shortage with the RTX 50 series. We're not sure where this ends, but be aware that prices will likely jump in the following months.

With that update wrapped, our budget builds have seen little in the way of changes. Fortunately, we've got some price drops, particularly for the Intel rig, with $30 off the motherboard, $25 on the CPU, and $10 on the graphics card. AMD prices remained stable, with one increase being the joint chassis, and a $2 saving on the RAM. That pushes our AMD build to $1,003—up $17—and pulls our Intel system down by $48 down to $893.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks XT View	$109
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$60
Mobo	GGigabyte B850 Eagle WiFi6E ATX **NEW**	$170
CPU	AMD Ryzen 5 9600X	$240
Cooler	be quiet! Pure Rock Slim 2	$25
GPU	ASRock Challenger SE OC A750 8GB **NEW**	$200
RAM	32GB (2x 16GB) Silicon Power XPOWER Pulse Gaming @ 6,400 C32	$86
SSD 1	500GB Crucial P3 Plus PCIe 4.0 M.2	$40
SSD 2	1TB Klevv CRAS C910 PCIe 4.0 M.2	$51
OS	Windows 11 Home 64-bit OEM	$22

Approximate Price: $1,003

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks XT View	$109
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$60
Mobo	MSI Pro Z790-S WiFi ATX	$110
CPU	Intel Core i5-14400F	$140
Cooler	be quiet! Pure Rock Slim 2	$25
GPU	XFX Speedster SWFT 210 Core RX 7600 8GB	$250
RAM	32GB (2x 16GB) Silicon Power Value Gaming @ 6,000 C30	$86
SSD 1	500GB Crucial P3 Plus PCIe 4.0 M.2	$40
SSD 2	1TB Klevv CRAS C910 PCIe 4.0 M.2	$51
OS	Windows 11 Home 64-bit OEM	$22

Approximate Price: $893

MID-RANGE

MOVING UP INTO our mid-range builds, and you can see across the board that we've actually enacted very little in the way of changes here, bar some major graphics card shake-ups.

Indeed, the market seems remarkably stable, all things considered, and we're seeing a number of price drops or just completely stable prices as a result. Hyte's Y40 chassis has bumped up in cost again; it's no longer available for the $90 it was there at, and is now sitting quite high at an impressive $130.

For our AMD system, prices have been quite stable. The Ryzen 7 9700X dropped by $24 on last month; the motherboard, cooler, and our primary MP600 Elite M.2 SSD have also remained stable. Thermaltake's Toughpower GF1 is at $95, and our memory kit and secondary SSD have seen a $3 and $2 price bump respectively.

The bigger problem is the graphics cards. Our AMD pick from last month, the ASRock Challenger OC RX 7700 XT, shot up by an astounding $168. We've done our best to find a suitable replacement in the Sapphire Pulse variant, but that's still pushing costs up by $76 or so on a fairly 'mid-range' GPU. Ouch. With that, our other price drops here have had little to no effect, with overall cost increasing to $1,549—up by $70 on last month.

Our Intel system has also seen a number of price increases, with the second largest jump being on that Core i7-14700F CPU, to $316 from $273. And then there's the GPU—last month's RTX 4070 Super had increased by $386. Sadly, try as we might, nowhere has an RTX 4070 Super for anything less than $900. We expect that even the price on our RTX 4060 Ti 8GB here isn't going to last, as the next available model up jumps by $100, making the average price for a 1080/1440p graphics card, $540. Nice.

Because we've dropped that cost there on the GPU, pricing has gone down overall by $94, bringing us to a nice round total of $1,497.

AMD INGREDIENTS

PART		PRICE
Case	Hyte Y40 ATX Mid Tower	$129
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	ASRock X870 Pro RS	$200
CPU	AMD Ryzen 7 9700X	$289
Cooler	Noctua U12S chromax.black	$90
GPU	Sapphire PULSE RX 7700 XT 12GB **NEW**	$465
RAM	32GB (2x 16GB) TeamGroup T-Create Expert @ 6,400 C32	$91
SSD 1	1TB Corsair MP600 Elite M.2 PCIe 4.0 SSD **NEW**	$75
SSD 2	2TB Silicon Power UD90 M.2 PCIe 4.0 SSD	$93
OS	Windows 11 Home 64-bit OEM	$22

Approximate Price: $1,549

INTEL INGREDIENTS

PART		PRICE
Case	Hyte Y40 ATX Mid Tower	$129
PSU	850W Thermaltake Toughpower GF1 2024 80+ Gold	$95
Mobo	ASRock Z790 PG Sonic ATX	$162
CPU	Intel Core i7-14700F	$316
Cooler	Razer Hanbo Chroma - 360mm AIO	$80
GPU	Asus Dual Evo OC RTX 4060 Ti 8GB **NEW**	$440
RAM	32GB (2x 16GB) Silicon Power Xpower Zenith Gaming @ 6,000 C30	$85
SSD 1	1TB Corsair MP600 Elite M.2 PCIe 4.0 SSD	$75
SSD 2	2TB Silicon Power UD90 M.2 PCIe 4.0 SSD	$93
OS	Windows 11 Home 64-bit OEM	$22

Approximate Price: $1,497

TURBO

GEOMETRIC FUTURE
05
GEOMETRIC FUTURE

IT FEELS LIKE A LOT OF manufacturers right now are aware of the issues facing the industry. Everywhere, we're seeing big changes. That's leading us to keep the majority of our system components resolutely identical to last month.

Our Geometric Future Model 5 chassis has bounced up by another $10, but generally for our AMD rig, we're seeing some major price drops: $31 on the CPU, $43 on the motherboard, $60 on the Crucial T700, and $12 on the Lexar NM790. Memory saw a slight uptick, by $10 or so, and the PSU bumped up by an outrageous $2, plus NZXT's Kraken 360 remained stuck at $180.

GPU time. Turbo has this kind of premise of being a 'no holds barred' PC, so downgrading graphics cards is a bit of a painful pill to swallow. Combine that with massive GPU shortages, and we'd have been dropping down to the RX 7800 XT, which feels wrong for a Turbo system. Equally, you can't get any RX 7900 GREs; there's one available, and it's priced the same as our XFX Speedster below. Almost all of the 7900 XTs right now are over $1,000. In fact, the white Speedster here is second cheapest to its black sibling that costs (for some reason) $100 more. If we'd stuck with the Phantom Gaming OC unit, prices would have bumped up by $380, so we'll take the $210 price hit on the chin. That leaves us with a total cost of $2,779—$86 up on last month.

The good news is it's $42 off on the motherboard, plus those SSD savings, and a smidge of dollar shifts here and there for everything else, but nothing to write home about—until we get to the GPU. There's no stock. If you want an RTX 4080 Super, you're talking $1,650 minimum, with prices climbing as high as $2,600. 5080s start at $1,750, and an RTX 4090 will set you back $3,300, with prices reaching $4,000 at some resellers (not eBay).

To counteract that, we've gone with MSI's RTX 5070 Ti. We've not personally tested this card yet, but it's the cheapest of the bunch, and early reports suggest it performs about 15-20 percent slower than an RTX 5080. So it's not all bad. Still, at this price, we'll quote *PC Gamer*: "You can knock 20 percent off the score [86%] if this one ends up a $900+ GPU."

AMD INGREDIENTS

PART		PRICE
Case	Geometric Future Model 5	$150
PSU	Super Flower LEADEX VII Platinum Pro 1000W - 80+ Platinum	$200
Mobo	Gigabyte X870E Aorus Elite WiFi7 ATX	$277
CPU	AMD Ryzen 9 9950X	$545
Cooler	NZXT Kraken 360 360mm AIO	$180
GPU	XFX Speedster MERC 310 RX 7900 XT 20GB White **NEW**	$900
RAM	64GB (2x 32GB) Patriot Viper Venom @ 6,400 C32	$150
SSD 1	2TB Crucial T700 PCIe 5.0 M.2	$210
SSD 2	2TB Lexar NM790 PCIe 4.0 M.2	$145
OS	Windows 11 Home 64-bit OEM	$22
Approximate Price:		**$2,779**

INTEL INGREDIENTS

PART		PRICE
Case	Geometric Future Model 5	$150
PSU	Super Flower LEADEX VII Platinum Pro 1000W - 80+ Platinum	$200
Mobo	Gigabyte Z890 Aorus Elite WiFi7	$218
CPU	Intel Core Ultra 9 285K	$600
Cooler	Asus ROG Strix LC II ARGB 360mm AIO	$155
GPU	MSI Ventus 3X OC RTX 5070 Ti 16GB **NEW**	$1,130
RAM	64GB (2x32GB) TeamGroup T-Create Expert @ 6,000 C34	$170
SSD 1	2TB Crucial T700 PCIe 5.0 M.2	$210
SSD 2	2TB Lexar NM790 PCIe 4.0 M.2	$145
OS	Windows 11 Home 64-bit OEM	$22
Approximate Price:		**$3,000**

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US LLC, 130 West 42nd Street, 7th Floor, New York, NY 10036. USA. Website: www.futureus.com. Future US LLC also publishes *Mac|Life*, and *PC Gamer*. Entire contents copyright 2025, Future US LLC. All rights reserved. Reproduction in whole or in part is prohibited. Future US LLC. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies.

PRODUCED IN THE UNITED STATES OF AMERICA.

UNEARTH THE FULL STORY BEHIND THE LEGEND OF ZELDA

Explore every game in the series and what makes each one so great. If you're a fan, you'll enjoy the trip through time. If you're a newcomer, we offer you this sage advice: "It's dangerous to go alone, take this!"

OLD!
retro GAMER PRESENTS
100% UNOFFICIAL
THE STORY OF
ZELDA
DISCOVER THE COMPLETE HISTORY OF HYRU
BRAND NEW
TEARS OF THE KINGDOM REVEALED
EVERY GAME IN THE SERIES EXPLORED
BEHIND-THE-SCENES ACCESS
TIMELINES EXPLAINED
Digital Edition
ON SALE NOW

FUTURE

Shop easily online at:
magazinesdirect.com

bit.ly/3xbOGXq

DIGITALSTORM

PC GAMER
EDITOR'S CHOICE
★★★★★
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM
Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

9001

9000